



User Guide

Amazon Q Business



Amazon Q Business: User Guide

Copyright © 2024 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

.....	X
What is Amazon Q Business?	1
Benefits of Amazon Q Business	1
Pricing and availability	2
Accessing Amazon Q Business	2
Related services	3
Are you a first-time Amazon Q Business user?	3
Getting started	5
How Amazon Q works	5
Admin workflow	5
User workflow	6
Amazon Q Business workflow	7
Key concepts	7
Document attributes and types	11
Understanding document attributes	11
Mapping document attributes	14
Supported document formats	18
Supported languages	37
Setting up	37
Initial AWS account setup	37
(Optional) Install the AWS CLI	39
(Optional) Set up the AWS SDKs	40
Consider AWS Regions and endpoints	40
Set up required permissions	41
Enable and configure an IAM Identity Center instance	41
IAM roles	42
IAM role for an Amazon Q Business application	43
IAM role for an Amazon Q web experience	45
IAM role for Amazon Q Business data source connectors	47
IAM role for Amazon Q Business plugins	58
IAM role for Lambda functions	60
IAM role for an Amazon Kendra retriever	65
Creating a sample application	67
Step 1: Create a sample application	67

Step 2: Add users and groups	70
Managing a sample application	74
Configuring an application	76
Creating an application	77
Managing applications	80
Creating a retriever	84
Amazon Q retriever	85
Amazon Kendra retriever	90
Connecting data sources	95
Upload documents	96
Amazon Kendra retriever	98
Amazon Q Business data sources	100
Adding users and groups	109
Data source connectors	114
Concepts	114
Source and endpoint metadata	115
Authorization	115
Authentication	116
Virtual private cloud	116
Web proxy	117
IAM role	117
Identity crawler	118
Sync scope	119
Sync mode	119
Sync run schedule	120
Field mappings	121
Configuration best practices	121
Supported connectors	122
AEM (Cloud)	124
AEM (Server)	173
Alfresco (Cloud)	222
Alfresco (Server)	245
Aurora (MySQL)	269
Aurora (PostgreSQL)	289
Amazon FSx (Windows)	309
Amazon RDS (Microsoft SQL Server)	326

Amazon RDS (MySQL)	346
Amazon RDS (Oracle)	367
Amazon RDS (PostgreSQL)	387
Amazon S3	407
Amazon Q custom connector	432
Amazon Q Web Crawler	436
Amazon WorkDocs	472
Box	490
Confluence (Cloud)	513
Confluence (Server/Data Center)	579
Dropbox	618
Drupal	642
GitHub (Cloud)	675
GitHub (Server)	709
Gmail	743
Google Drive	770
IBM DB2	800
Jira	820
Microsoft Exchange	859
Microsoft OneDrive	889
Microsoft SharePoint (Online)	911
Microsoft SharePoint (Server)	966
Microsoft SQL Server	1014
Microsoft Teams	1035
Microsoft Yammer	1077
MySQL	1104
Oracle Database	1124
PostgreSQL	1144
Quip	1164
Salesforce Online	1188
ServiceNow Online	1297
Slack	1344
Zendesk	1378
Understanding User Store	1418
Principal mapping	1418
How the User Store works	1421

Using Amazon VPC	116
Configuring Amazon VPC	1423
Connecting to Amazon VPC	1425
Using Amazon VPC with Amazon S3	1427
Connecting to a database	1431
Troubleshooting VPC connection issues	1433
Troubleshooting data source connectors	1435
My documents were not indexed	1436
My synchronization job failed	1436
My synchronization job is incomplete	1437
My synchronization job succeeded but there are no indexed documents	1437
I am running into file format issues while syncing my data source	1438
I am getting an AccessDenied When Using SSL Certificate File error message	1438
Enhancing an application	1439
Document enrichment	1439
How document enrichment works	1440
Using basic operations	1442
Using Lambda functions	1451
Plugins	1472
Jira plugin	1473
Salesforce plugin	1475
ServiceNow plugin	1478
Zendesk plugin	1481
Using plugins	1484
Managing plugins	1487
Admin controls and guardrails	1492
Key terms	1492
Using global controls	1493
Using topic-level controls	1497
Managing admin controls and guardrails	1501
Relevance tuning	1503
Understanding boosting	1504
Boosting types	1505
Configuring document attributes for boosting	1507
Enabling document attributes for search	1514
Customizing a web experience	1518

Customize web experience	1518
Managing web experiences	1520
Deleting a web experience	1520
Getting properties of a web experience	1521
Listing web experiences	1522
Using a web experience	1523
Prompts	1523
Engage with contextual responses	1524
Analyze content	1524
Perform actions on your behalf	1524
Review source citations	1524
Upload files and chat	1524
Copy responses	1525
Provide feedback	1525
Conversation management	1525
Conversation settings	1526
Retrieve conversation IDs	1526
Amazon Q Business features	1532
Filtering using document attributes	1532
Source attribution with citations	1533
Upload files and chat	1534
Quick prompts	1534
Legacy identity management	1536
How it works	1536
Create an Amazon Q application	1537
Previewing a web experience	1540
Preview and customize web experience	1541
Testing Amazon Q Business web experience functions	1543
Managing web experiences	1544
Creating a retriever	1549
Amazon Q retriever	1550
Amazon Kendra retriever	1555
Connecting data sources	1560
Upload documents	1561
Amazon Kendra retriever	1563
Amazon Q Business data sources	1565

Deploying a web experience	1574
Integration process overview	1575
Key IdP integration concepts	1578
Deploying instructions	1581
Troubleshooting IdP integration	1625
Tagging resources	1629
Using tags	1629
Tag restrictions	1630
Security	1631
Data protection	1632
Data encryption	1633
Key management	1643
Identity and access management	1644
Audience	1644
Authenticating with identities	1645
Managing access using policies	1648
How Amazon Q Business works with IAM	1650
Identity-based policy examples	1659
Troubleshooting	1666
Compliance validation	1667
Resilience	1669
Infrastructure security	1669
Cross-service confused deputy prevention	1669
Configuration and vulnerability analysis	1671
Security best practices	1671
Apply principle of least privilege	1671
Role-based access control (RBAC) permissions	1672
Monitoring	1673
CloudTrail logs	1673
Amazon Q information in CloudTrail	1673
Control plane events in CloudTrail	1674
Data plane events in CloudTrail	1676
Amazon Q management events in CloudTrail	1678
Understanding Amazon Q log file entries	1678
Service quotas	1681
Supported Regions	1681

Quotas	1681
API reference	1683
Overview of Amazon Q API operations	1683
Creating an application	1684
Creating an index	1684
Creating a retriever	1685
Connecting data sources	1685
Upload documents directly	1685
Creating and customizing a web experience	1686
Chat and conversation management	1686
User and group management	1686
Amazon Q plugins	1687
Admin controls and guardrails	1687
User feedback	1687
Document history	1688

Amazon Q is in preview release and is subject to change.

What is Amazon Q Business?

Powered by Amazon Bedrock: AWS implements [automated abuse detection](#). Because Amazon Q is built on Amazon Bedrock, users can take full advantage of the controls implemented in Amazon Bedrock to enforce safety, security, and the responsible use of artificial intelligence (AI).

Amazon Q Business is a fully managed, generative-AI powered assistant that you can configure to answer questions, provide summaries, generate content, and complete tasks based on your enterprise data. Amazon Q provides immediate and relevant information to employees, and helps streamline tasks and accelerate problem solving.

Amazon Q integrates with services like [Amazon Kendra](#) and [other supported data sources](#) such as [Amazon S3](#), [Microsoft SharePoint](#), and [Salesforce](#).

[What is Amazon Q Business?](#)

Topics

- [Benefits of Amazon Q Business](#)
- [Pricing and availability](#)
- [Accessing Amazon Q Business](#)
- [Related services](#)
- [Are you a first-time Amazon Q Business user?](#)

Benefits of Amazon Q Business

Some of the benefits of Amazon Q Business include:

Accurate and comprehensive answers

Amazon Q generates comprehensive responses to natural language queries from users by analyzing information across all enterprise content that it has access to. It can avoid incorrect statements by confining its generated responses to existing enterprise data. Amazon Q also provides citations to the sources that it used to generate its response.

Simple to deploy and manage

Amazon Q undertakes the complex task of developing and managing machine learning infrastructure and models so that you can build your chat solution quickly. Amazon Q connects to your data and ingests it for processing using its pre-built connectors, document retrievers, document upload capabilities.

Configurable and customizable

Amazon Q provides you with the flexibility of choosing what sources should be used to respond to user queries. You can control whether the responses should only use your enterprise data, or use both enterprise data and model knowledge.

Data and application security

Amazon Q supports access control for your data so that the right users can access the right content. Its responses to questions are based on the content that your end user has permissions to access. You can integrate with IAM Identity Center to manage user access.

Broad connectivity

Amazon Q offers out-of-the-box connections to multiple supported data sources. For more information about data source connectors supported by Amazon Q, see [Amazon Q data source connectors](#).

Pricing and availability

Amazon Q is currently in preview. For more information on pricing during the preview, see [Amazon Q pricing](#).

For a list of regions where Amazon Q is currently available, see [Supported regions](#).

During Preview, an Amazon Q application supports only 50 end users. If you need more capacity, contact [Support](#).

Accessing Amazon Q Business

You can access Amazon Q Business in the following ways in the AWS Regions that it's available in:

[AWS Management Console](#)

You can use the AWS Management Console—a browser-based interface to interact with AWS services—to access the Amazon Q console and resources. You can perform most Amazon Q tasks using the Amazon Q console.

[Amazon Q Business API](#)

To access Amazon Q Business programmatically, you can use the Amazon Q API. For more information, see the [Amazon Q API Reference](#).

[AWS Command Line Interface](#)

The AWS Command Line Interface (AWS CLI) is an open source tool. You can use the AWS CLI to interact with AWS services using commands in your command line shell. If you want to build task-based scripts, using the command line can be faster and more convenient than using the console.

[SDKs](#)

AWS SDKs provide language APIs for AWS services to use programmatically.

Related services

The following are some of the other AWS services that Amazon Q Business integrates with:

[Amazon Kendra](#)

Amazon Kendra is an intelligent search service that uses natural language processing and machine learning algorithms to return specific answers from your data for end user queries. If you're already an Amazon Kendra user, you can use Amazon Kendra as a data retriever for your Amazon Q web application.

[Amazon S3](#)

Amazon S3 is an object storage service. If you're an Amazon S3 user, you can use Amazon S3 as a data source for your Amazon Q application.

Are you a first-time Amazon Q Business user?

If you're a first-time user of Amazon Q Business, we recommend that you read the following sections in order:

[How it works](#)

Introduces Amazon Q components and describes how they work to create your Retrieval Augmented Generation (RAG) solution.

[Key concepts](#)

Explains key concepts and important Amazon Q terminology.

[Setting up](#)

Explains key concepts and important Amazon Q terminology and outlines how to set up Amazon Q so that you can begin creating your Amazon Q application and web experience.

[Creating an application](#)

Explains how to create the Amazon Q application that powers your Amazon Q web experience.

[Configuring Amazon Q Business data source connectors](#)

Configuration information for specific connectors to use with your Amazon Q web experience.

Getting started

To start using Amazon Q Business, set up an AWS account and create the necessary AWS Identity and Access Management (IAM) users and roles. To use the AWS Command Line Interface (AWS CLI) or the AWS SDKs, you must install and configure them. After learning about Amazon Q concepts and setting up, you are ready to begin creating your application.

Topics

- [How Amazon Q Business works](#)
- [Key concepts of Amazon Q Business](#)
- [Document attributes and types in Amazon Q Business](#)
- [Supported languages for Amazon Q Business](#)
- [Setting up for Amazon Q Business](#)
- [IAM roles for Amazon Q Business](#)

How Amazon Q Business works

With Amazon Q Business, you can build an interactive chat application for your organization's end users, using a combination of your enterprise data and large language model knowledge, or enterprise data only. The following sections outline how Amazon Q works.

Topics

- [Admin workflow](#)
- [User workflow](#)
- [Amazon Q Business workflow](#)

Admin workflow

As an admin user, you create and configure an Amazon Q web experience by completing the following steps:

1. [Configuring an IAM Identity Center instance](#) for your Amazon Q application. Your IAM Identity Center instance must be created in the same region as your Amazon Q application.

2. (Optional) [Creating a sample Amazon Q application](#) to test how Amazon Q Business works before [creating a fully-configured application](#).
3. [Creating the Amazon Q application](#) that powers your web experience and connecting it to IAM Identity Center.
4. [Choosing a retriever](#) for the application.
5. [Connecting your data sources](#) to—or uploading data into—the application.
6. [Adding groups and users](#) who will access the Amazon Q web experience. This includes adding at least one user to your application.
7. [Enhancing the web experience](#) by configuring admin-level controls, relevance tuning, and the end user chat experience. For more information, see [Enhancing an Amazon Q application](#) and [Amazon Q features](#).
8. Optionally, [customizing your web experience](#) to test how it looks for your end users. In this step, you add a title and subtitle for your web experience, and a welcome message for your end users. You can't chat with—or test—the application in customize mode.

During Preview, an Amazon Q application supports only 50 end users. If you need more capacity, contact [Support](#).

Note

If you're looking for information on legacy SAML 2.0 identity management systems supported by Amazon Q Business, see [Legacy identity management](#).

User workflow

If you're an end user using your organization's Amazon Q Business web experience, you perform the following steps:

1. Navigate to your organization's Amazon Q web experience URL, and sign in with your credentials.
2. Start chatting and ask questions of your organization's Amazon Q web experience. For a list of web experience capabilities, see [Using an Amazon Q web experience](#).
3. Sometimes your question requires information that's beyond the scope of your enterprise data. Then, Amazon Q responds that it couldn't find an answer in your documents, unless your admin has allowed Amazon Q to [generate responses using model knowledge](#).

4. Additionally, you can ask Amazon Q to complete [any supported follow-up tasks](#)—like creating a Jira ticket—that your admin has configured.

Amazon Q stores conversation history for 30 days and maintains conversation context after a conversation ends. Conversations can be resumed from where you left off within this 30-day period.

Amazon Q Business workflow

In response to an end user query during a web experience chat, Amazon Q does the following:

1. Uses the retriever chosen by the admin to select and retrieve documents that are relevant to the query, following authorization and access control.
2. Generates a response to the user query using either a combination of retrieved enterprise data and model knowledge, or only enterprise data, depending on admin configuration.
3. Returns the generated response to the end user. Amazon Q assigns a unique message ID to each answer for tracking purposes.

Key concepts of Amazon Q Business

This section describes the key concepts and terms related to Amazon Q Business.

Retrieval Augmented Generation

Retrieval Augmented Generation (RAG) is a natural language processing (NLP) technique. Using RAG, generative artificial intelligence (generative AI) is conditioned on specific documents that are retrieved from a dataset. Amazon Q has a built-in RAG system. A RAG model has the following two components:

- A *retrieval* component retrieves relevant documents for the user query.
- A *generation* component takes the query and the retrieved documents and then generates an answer to the query using a large language model.

Large language model

A large language model (LLM) is a language-based, machine learning model that's tuned to a large number (billions) of parameters or and trained on a large corpus of documents.

Retriever

A retriever pulls data from an index in real time during a conversation. Amazon Q supports a native index retriever and also a Amazon Kendra index retriever.

Index

An index is a corpus of documents. Amazon Q supports its own index where you can add and sync documents. An index has fields that you can map your document attributes to enhance your end user's chat experience. Amazon Q creates an index for you when it creates your Amazon Q native retriever.

You can also use an Amazon Kendra index as a retriever for your generative AI application.

Data source

A data source is a document repository.

Data source connector

A data source connector can crawl and synchronize a data source with an Amazon Q index at customizable intervals. Amazon Q supports multiple connectors so that you can build your generative AI solution with minimal configuring. For a list of Amazon Q supported connectors, see [Supported connectors](#). For an overview of Amazon Q connector features, see [Amazon Q data source connector features](#).

IAM Identity Center

You can manage user access to your Amazon Q Business application using IAM Identity Center as your AWS gateway to the identity provider of your choice. For more information on creating an Amazon Q Business application integrated with IAM Identity Center see [Configuring an Amazon Q Business application](#). For more information about using IAM Identity Center to manage access to applications, see [Manage access to applications](#) in the IAM Identity Center User Guide.

Document

In Amazon Q, a document is a unit of data. Specific document formats supported include .csv, .docx, HTML, JSON, .pdf, plaintext, .ppt, .pptx, .rtf, and .xlsx. Amazon Q supports both structured and unstructured text. For more information, see [Supported document types](#).

Application

An Amazon Q application is the primary resource that you use to create a chat solution. To create the application, you can use either the Amazon Q console or [Amazon Q API](#) actions.

Web experience

An Amazon Q web experience is the chat interface that you create using your Amazon Q application. Then, your end users can chat with your organization's Amazon Q web experience. You can configure and customize your Amazon Q web experience using either the Amazon Q console or the Amazon Q API.

Guardrails and chat controls

An Amazon Q feature that lets you define global controls and topic-level controls for your application. Using this feature, you can control what sources your application will use to generate responses from, and also control what topics it will respond to and how. For more information, see [Guardrails](#).

Plugins

Amazon Q includes a plugins feature that you can use to interact with third-party services such as Jira and Salesforce. With the plugins feature, you can perform actions specific to that service (like creating a ticket) from within your Amazon Q web experience chat. For more information, see [Plugins](#).

Quick prompts

The Amazon Q quick prompts feature helps with end user discoverability of the web experience chat features. Use this feature to prompt your end user to engage with their web experience chat in specific ways. For example, you can show the available [configured plugins](#) or inform users that they can choose to summarize their chat.

Document attributes

Document attributes are structural metadata associated with documents, such as document title, document type, and date and time created. Amazon Q extracts document attributes during the document ingestion process to provide customizable chat and data manipulation capabilities for your application. Amazon Q offers reserved document attributes that you can use. Or, you can create custom attributes. For more information, see [Document attributes](#), [Filtering using document attributes](#), [Boosting using document attributes](#), and [Custom document enrichment](#).

Filtering using document attributes

Filtering using document attributes is an Amazon Q feature that you can use to filter your Amazon Q chat responses for your end user. For example, if you have a document attribute associated with a data source type, you can use the attribute to mandate that chat responses

only be generated from a specific data source. For more information, see [Filtering using document attributes](#).

Relevance tuning

You can choose to use document attributes to boost and tune the relevance of chat responses for end users from specific content. For example, if you have a document attribute associated with document creation or update date, you use these attributes to boost chat responses from more recently created or updated documents. For more information, see [Relevance tuning](#).

Document enrichment

Document enrichment is an Amazon Q feature that you can use to manipulate your document content and document attributes. You can use document enrichment to perform optical character recognition (OCR) or translation. Document enrichment uses basic and Lambda operations. For more information see, [Document attributes and types](#) and [Document enrichment](#).

Field mappings

An Amazon Q index has fields that help you structure data to aid the retrieval process. You can map index fields to your [document attributes](#) when you add documents directly to an index, or use a data source connector.

User Store

User Store is an Amazon Q data source connector feature that streamlines user and group management across all the data sources attached to your application. For more information about how this feature works and implementation details, see [Understanding User Store](#).

Index units

When you use an Amazon Q native retriever for your application, you must provision data storage capacity for your index. You can provision between 1–50 units for your index. Each unit is equal to 20,000 documents or 200 MB, whichever comes first. For more information, see [Amazon Q pricing](#).

Tags

Manage your Amazon Q applications and data sources by assigning tags or labels. You can use tags to categorize your Amazon Q resources in various ways. For example, categorize by purpose, owner, or application, or any combination. Each tag consists of a key and a value, both of which you define. For more information, see [Tags](#).

Foundation model

A foundation model (FM) is a broad, function-based machine learning model (not specific to language systems). An FM is tuned to a large number (billions) of parameters and is trained on a large corpus of documents.

Hallucination

A hallucination, in the machine learning context, is a confident response by an AI application that isn't justified by its training data. Think of a hallucination as instances where the response doesn't make sense in the context of the prompt, or when the responses are out of scope with the documents provided. Amazon Q offers you the ability to minimize hallucinations by allowing your retrieval system to [generate responses only from your existing enterprise data](#).

Document attributes and types in Amazon Q Business

This section outlines what document attributes are, how they work in Amazon Q Business, and what they can help you do for your chat solution. This section also lists the document types supported by Amazon Q.

Topics

- [Understanding document attributes in Amazon Q Business](#)
- [Mapping document attributes in Amazon Q Business](#)
- [Supported document formats in Amazon Q Business](#)

Understanding document attributes in Amazon Q Business

Every document has structural attributes—or metadata—attached to it. Document attributes can include information such as document title, document author, time created, time updated, and document type.

You can map document attributes to fields in your Amazon Q Business index. Once mapped to document attributes, these index fields can be used by admin to boost results from specific sources, or by end users to filter and scope their chat results to specific data.

Note

During Preview, filtering using document attributes in chat is only supported through the API. Boosting search results using document attributes is supported on both the console and the API.

You can use document attributes to prepare your data for—and customize and control—end user chat. To learn more, see [Filtering using metadata](#), [Document enrichment in Amazon Q](#), and [Relevance tuning](#).

Topics

- [Types of document attributes](#)
- [Mapped document attributes](#)
- [Document attribute data types](#)

Types of document attributes

Amazon Q Business supports two types of document attributes: reserved and custom.

Reserved or default document attributes are provided by Amazon Q to map commonly occurring document attributes to index fields. Custom attributes, on the other hand, can be used to map document attributes unique to your content to index fields.

Both reserved and custom document attributes can be used to customize end user chat experience.

The following section outlines the available document attributes.

Topics

- [Reserved document attributes](#)
- [Custom document attributes](#)

Reserved document attributes

Amazon Q Business offers the following reserved document attributes or index fields that you can map your metadata to:

- `_authors` – A list of one or more authors responsible for the content of the document.

- `_category` – A category that places a document in a specific group.
- `_created_at` – The date and time in ISO 8601 format that the document was created. For example, 2012-03-25T12:30:10+01:00 is the ISO 8601 date-time format for March 25, 2012 at 12:30 PM (plus 10 seconds) in Central European Time.
- `_data_source_id` – The identifier of the data source that contains the document.
- `_document_body` – The content of the document.
- `_document_id` – A unique identifier for the document.
- `_document_title` – The title of the document.
- `_file_type` – The file type of the document, such as .pdf or .docx.
- `_last_updated_at` – The date and time in ISO 8601 format that the document was last updated. For example, 2012-03-25T12:30:10+01:00 is the ISO 8601 date-time format for March 25, 2012 at 12:30 PM (plus 10 seconds) in Central European Time.
- `_source_uri` – The URI where the document is available. For example, the URI of the document on a company website.
- `_version` – An identifier for the specific version of a document.
- `_view_count` – The number of times that the document has been viewed.
- `_language_code` (String) – The code for a language that applies to the document. This defaults to English if you don't specify a language.

Custom document attributes

You can also create custom attributes based on your own enterprise data. Then, you can map the custom attributes to custom index fields that you create for a more tailored end user chat experience.

For example, you can create a custom field or attribute called "Department" with the values of "HR", "Sales", and "Manufacturing". Then, you can use these fields or attributes to allow your end users to filter their chat results to documents in the "HR" department, or restrict response generation to specific data stores.

You can create up to 50 custom fields or attributes.

Important

Once created, you can't delete or rename any attributes.

Mapped document attributes

When a document attribute—reserved or custom—is mapped to an index field, you can choose how the field will be used during chat. You can currently configure index fields to perform the following action:

- **Search** – Allows end users the ability to search data with the specified attributes.

Document attribute data types

Document attributes—reserved or custom—can only be the data types that are shown in the following table. Additionally, document attributes can be used to perform the operations outlined.

Data type	Searchable	Filterable	Boostable		
Date	No	Yes	Yes		
Number	No	Yes	Yes		
String	Yes	Yes	Yes		
String list	Yes	Yes	Yes		

For more information on filtering and boosting using document attributes, see [Filtering using document-attributes](#) and [Boosting using document attributes](#).

Note

You can't change an index field type after it has been created.

Mapping document attributes in Amazon Q Business

An Amazon Q Business index has field you can map your document attributes to. Once mapped to document attributes, these index fields can be used by admin to boost results from specific sources, or by end users to filter and scope their chat results to specific data.

Mapping document attributes from your documents to index fields is a multi-step process that depends on the document upload method you use.

Note

During Preview, filtering using document attributes in chat is only supported through the API. Boosting search results using document attributes is supported on both the console and the API.

Topics

- [Mapping document attributes directly to index fields](#)
- [Mapping data source document attributes to index fields](#)
- [Ingesting attributes using the BatchPutDocument API operation](#)

Mapping document attributes directly to index fields

When you use the API, you must first map your document attributes to index fields before you can use them for filtering in chat. You use the following process to map document attributes to your index field:

1. You create an index by calling the [CreateIndex](#) API operation.
2. Then, you create index fields using the [UpdateIndex](#) operation. You use this method to map both reserved and custom document attributes to index fields.
3. Optionally, you can test and view the index fields that you've added by using the [GetIndex](#) operation.
4. Then, when you use the [BatchPutDocument](#) operation to ingest documents into your index, Amazon Q extracts your reserved or custom document attributes and maps them to the index fields that you have already created.

After you map document attributes directly to index fields using the API, you can select specific attributes for your end user to use for filtering chat responses. With the `UpdateIndex` API operation, you add custom fields or attributes using the `documentAttributeConfigurations` parameter.

The following JSON example uses `documentAttributeConfigurations` to add a field called "Department" to the index.

```
"DocumentmetadataConfigurationUpdates": [
```

```
{
  "Name": "Department",
  "Type": "STRING_VALUE"
}
```

Mapping data source document attributes to index fields

If you use an Amazon Q Business data source connector, you can map default document attributes attached to documents in your data source to fields in your Amazon Q index. You can use these document attributes to help your end user filter and scope chat responses.

Important

During Preview, filtering using data source document attributes in chat is only supported through the API.

Each data source connector is designed to crawl the default document attributes in your data source automatically. For example, if you have a field in your data source named `dept` that contains department information for a document, you can map it to an index field named `Department`. You can't change or customize default data source attributes that are mapped to an index.

You can also map any Amazon Q reserved fields such as `_created_at`. If your data source has a field named `creation_date`, you can map this field to the equivalent Amazon Q reserved field named `_created_at`.

You can also choose to add custom document attributes and map them to custom fields that you create in your index. Most data sources support field mappings and follow a specific configuration format, except Amazon S3 and database data sources:

- If you store your documents in an Amazon S3 bucket or Amazon S3 data source, you can either use the console to specify field mappings or specify fields [using a JSON metadata file](#).

When you use an Amazon S3 bucket as a data source for your index, you use companion metadata files to add metadata to the documents. You place the metadata JSON files in a directory structure that is parallel to your documents. For more information, see [S3 document metadata](#).

You specify custom fields or attributes in the `Attributes` JSON structure. You can create up to 50 custom fields or attributes. The following example uses `Attributes` to define three custom fields or attributes and one reserved field.

```
"Attributes": {
  "brand": "Amazon Basics",
  "price": 1595,
  "_category": "sports",
  "subcategories": ["outdoors", "electronics"]
}
```

- For database data sources, if the name of the database column matches the name of a reserved field, the field and column are mapped automatically.

If you use the console, you select default field mappings or create custom mappings when you configure your connector. On the console, if a default field or a default field property can't be edited, it will appear grayed out.

If you use the API, you use the `configuration` parameter of the [CreateDataSource](#) API operation to map default document attributes in your data source to index fields.

If you want to map custom document attributes in your data source to Amazon Q index fields, use the `DocumentAttribute` parameter of the [UpdateIndex](#) operation to first create the custom field matching the custom document attribute. By doing so, you can specify and map your reserved or custom data source document attribute to a reserved or custom index field.

Ingesting attributes using the `BatchPutDocument` API operation

When you use the [BatchPutDocument](#) API operation to add a document to your index, you can specify document attributes—both reserved and custom—as part of `Attributes`. You can add multiple fields or attributes when you call the API operation. You can create up to 50 custom fields or attributes. The following example is a custom field or attribute that adds "Department" to a document.

```
"Attributes":
{
  "Department": "HR",
  "_category": "Vacation policy"
}
```

Supported document formats in Amazon Q Business

You can add documents to an Amazon Q Business application in three ways:

- [Using direct document upload](#) – If you use an Amazon Q retriever, you can directly upload documents into your application.
- [Using data source connectors](#) – You can add documents to an Amazon Q application using both the console and the API.
- [Using Upload files and chat](#) – As an end user using an Amazon Q web experience, you can directly upload up to 5 files during a conversation.

When you add documents to an Amazon Q application (directly or through datasource connectors) using the console or the API, Amazon Q extracts document content and internally parses these to optimize chat responses. The maximum file size of a single document must be 50 MB or less. The maximum amount of text that can be extracted from a single document is 5 MB.

When you upload documents using the documents directly into chat using the [Upload files and chat](#) feature, the size of each file you upload must be 10 MB or less. The total parsed content for all files combined have to be under 30,000 tokens or 20,000 words. 1 word corresponds roughly to 1.5 tokens.

Additionally, if you're uploading Comma Separated Values (CSV) or Microsoft Excel (XLS and XLSX) documents directly into chat, Amazon Q performs best for tables with approximately 4 columns and 10 rows. Files indexed by an Amazon Q data source connector or uploaded directly have no such restrictions.

Along with specific formats like PDF, Word, for example, each enterprise data source also has different entities that it considers documents. To learn about supported entity types for each data source, see [What is a document?](#)

Topics

- [Supported document types](#)
- [What is a document?](#)

Supported document types

The following table shows the document formats that Amazon Q Business supports.

Document format	How document is treated		
Portable Document Format (PDF)	Converted to HTML, then plain text is extracted. Scanned PDFs aren't supported as they are images.		
HyperText Markup Language (HTML)	HTML tags are filtered out to extract plain text. Content must be between the main HTML start and closing tags (<HTML>content</HTML>).		
Extensible Markup Language (XML)	XML tags are filtered out and plain text is extracted.		
Extensible Stylesheet Language Transformations (XSLT)	Tags are filtered out to extract plain text.		
Markdown (MD)	Content is extracted as plain text with Markdown syntax retained.		
Comma Separated Values (CSV)	Content is extracted as plain text from each cell, with a single file treated as a single document result. Amazon Q		

Document format	How document is treated		
	<p>doesn't support analytics questions for CSVs; it supports only qualitative questions.</p>		
<p>Microsoft Excel (XLS and XLSX)</p>	<p>Content is extracted as plain text from each cell, with a single row treated as a single document result. Amazon Q doesn't support analytics questions for Excel files; it supports only qualitative questions.</p>		
<p>JavaScript Object Notation (JSON)</p>	<p>Content is extracted as plain text with JSON syntax retained.</p>		
<p>Rich Text Format (RTF)</p>	<p>RTF syntax is filtered out to extract plain text content.</p>		
<p>Microsoft PowerPoint (PPT, PPTX)</p>	<p>Only plain text content is extracted from PowerPoint slides for ingestion . Images and other content aren't extracted.</p>		

Document format	How document is treated		
Microsoft Word (DOCX)	Only plain text content is extracted from Word pages for ingestion. Images and other content aren't extracted.		
Plain text (TXT)	All text in the text document is extracted.		

What is a document?

When you directly add files to Amazon Q Business using the [Using direct document upload](#) or the [Upload files and chat](#) feature, it considers each file you add a document. When you connect Amazon Q to a data source, what Amazon Q considers—and crawls—as a document varies by connector.

The following table outlines what each connector crawls as a document.

Data source connector	Supports crawling	Document definition	
Adobe Experience Manager (Cloud and Server)	<ul style="list-style-type: none"> Assets Pages 	<ul style="list-style-type: none"> Each Asset is considered a single document. Each Page is considered a single document. 	
Alfresco (Cloud and Server)	<ul style="list-style-type: none"> Files Comments 	<ul style="list-style-type: none"> Each File is considered a single document. 	

Data source connector	Supports crawling	Document definition	
		<ul style="list-style-type: none"> Each Comment is considered a single document. 	
Amazon FSx (Windows)	Files	Each File is considered a single document.	
Amazon S3	Objects	<p>Each Object is considered a single document.</p> <p>Any <i>object-name.metadata.json</i> file and access control list (ACL) file is considered metadata for the object it is associated with and not treated as a separate document.</p>	
Amazon Q Web Crawler	<ul style="list-style-type: none"> Web pages Attachments 	<ul style="list-style-type: none"> Each Web page is considered a single document. Each Attachment is considered a single document. 	
Amazon WorkDocs	<ul style="list-style-type: none"> Files Comments 	<ul style="list-style-type: none"> Each File is considered a single document. Each Comment is considered a single document. 	

Data source connector	Supports crawling	Document definition	
Box	<ul style="list-style-type: none"> • Files • Tasks • Comments • Weblinks 	<ul style="list-style-type: none"> • Each File is considered a single document. • Each Task is considered a single document. • Each Comment is considered a single document. • Each Weblink is considered a single document. 	
Confluence (Cloud and Server)	<ul style="list-style-type: none"> • Spaces • Pages • Blogs • Comments • Attachments 	<ul style="list-style-type: none"> • Each Space is considered a single document. • Each Page is considered a single document. • Each Blog is considered a single document. • Each Comment is considered a single document. • Each Attachment is considered a single document. 	

Data source connector	Supports crawling	Document definition	
<p>Database data sources</p> <ul style="list-style-type: none"> • Aurora (MySQL) • Aurora (PostgreSQL) • Amazon RDS (Microsoft SQL Server) • Amazon RDS (MySQL) • Amazon RDS (Oracle) • Amazon RDS (PostgreSQL) • IBM DB2 • PostgreSQL • Microsoft SQL Server • MySQL • Oracle Database 	<ul style="list-style-type: none"> • Table data in a single database • View data in a single database 	<p>Each row in a table and view is considered a single document.</p>	

Data source connector	Supports crawling	Document definition	
Dropbox	<ul style="list-style-type: none">• Files• Papers• Paper templates• Shortcuts	<ul style="list-style-type: none">• Each File is considered a single document.• Each Paper is considered a single document.• Each Paper template is considered a single document.• Each Shortcut is considered a single document.	

Data source connector	Supports crawling	Document definition	
Drupal	<ul style="list-style-type: none"> • Articles • Basic pages • Basic blocks • Custom content • Custom blocks • Comments on articles, basic pages, basic blocks, custom content, and custom blocks • Attachments in articles, basic pages, basic blocks, custom content, and custom blocks 	<ul style="list-style-type: none"> • Each Article is considered a single document. • Each Basic page is considered a single document. • Each Basic block is considered a single document. • Each Custom content is considered a single document. • Each Custom block is considered a single document. • Each Comment on an article, a basic page, a basic block, any custom content, and a custom block is considered a document. • Each Attachment in an article, a basic page, a basic block, any custom content, and a custom block is considered a document. 	

Data source connector	Supports crawling	Document definition	
GitHub (Cloud and Server)	<ul style="list-style-type: none"> • Repositories • Repository commits • Issues • Issue attachments • Issue comments • Pull request documents • Pull request comments • Pull request attachments 	<ul style="list-style-type: none"> • Each Repository is considered a single document. • Each Repository commit is considered a single document. • Each Issue is considered a single document. • Each Issue attachment is considered a single document. • Each Issue comment is considered a single document. • Each Pull request is considered a single document. • Each Pull request comment is considered a single document. • Each Pull request attachment is considered a single document. 	

Data source connector	Supports crawling	Document definition	
Gmail	<ul style="list-style-type: none"> • Emails • Email attachments 	<ul style="list-style-type: none"> • Each Email is considered a single document. • Each Email attachment is considered a single document. 	
Google Drive	<ul style="list-style-type: none"> • Files • Comments 	<ul style="list-style-type: none"> • Each File is considered a single document. • Each Comment is considered a single document. 	
Jira	<ul style="list-style-type: none"> • Projects • Issues • Comments • Attachments • Worklog 	<ul style="list-style-type: none"> • Each Project is considered a single document. • Each Issue is considered a single document. • Each Comment is considered a single document. • Each Attachment is considered a single document. • Each Worklog is considered a single document. 	

Data source connector	Supports crawling	Document definition	
Microsoft Exchange	<ul style="list-style-type: none"> • Emails • Attachments • Calendar • Contacts • Notes • OneNotes 	<ul style="list-style-type: none"> • Each Email is considered a single document. • Each Attachment is considered a single document. • Each Calendar is considered a single document. • Each Contact is considered a single document. • Each Note is considered a single document. • Each page in OneNotes is considered a single document. 	
Microsoft OneDrive	<ul style="list-style-type: none"> • Files • OneNotes 	<ul style="list-style-type: none"> • Each File is considered a single document. • Each page in OneNotes is considered a single document. 	

Data source connector	Supports crawling	Document definition	
Microsoft SharePoint (Online and Server)	<ul style="list-style-type: none"> • Events • Pages • Files • Links • File attachments • Comments • OneNotes 	<ul style="list-style-type: none"> • Each Event is considered a single document. • Each Page is considered a single document. • Each File is considered a single document. • Each Link is considered a single document. • Each File attachment is considered a single document. • Each Comment is considered a single document. • Each page in OneNotes is considered a single document. 	

Data source connector	Supports crawling	Document definition	
Microsoft Teams	<ul style="list-style-type: none"> • Chat messages • Chat attachments • Channel posts • Channel wikis • Channel attachments • Meeting chats • Meeting files • Meeting notes • Calendar meetings • OneNotes 	<ul style="list-style-type: none"> • Each Chat message is considered a single document. • Each Chat attachment is considered a single document. • Each Channel post is considered a single document. • Each Channel wiki is considered a single document. • Each Channel attachment is considered a single document. • Each Meeting chat is considered a single document. • Each Meeting file is considered a single document. • Each Meeting note is considered a single document. • Each Calendar meeting is considered a single document. • Each page in OneNotes is 	

Data source connector	Supports crawling	Document definition	
		considered a single document.	
Microsoft Yammer	<ul style="list-style-type: none"> • Communities • Attachments • Messages • Users 	<ul style="list-style-type: none"> • Each Community is considered a single document. • Each Attachment is considered a single document. • Each Message and community post is considered a single document. • Each User is considered a single document. 	
Quip	<ul style="list-style-type: none"> • Files • Messages • Threads 	<ul style="list-style-type: none"> • Each File is considered a single document. • Each Comment is considered a single document. • Each file and message posted in a Thread is considered a single document. 	

Data source connector	Supports crawling	Document definition	
Salesforce	<ul style="list-style-type: none"> • Accounts • Contacts • Campaigns • Contracts • Cases • Partners • Opportunities • Groups • Leads • Users • Tasks • Ideas • Profiles • Solutions • Chatters • Documents • Custom entities • Knowledge articles 	<ul style="list-style-type: none"> • Each Account is considered a single document. • Each Contact is considered a single document. • Each Campaign is considered a single document. • Each Contract is considered a single document. • Each Case is considered a single document. • Each Partner is considered a single document. • Each Opportunity is considered a single document. • Each Group is considered a single document. • Each Lead is considered a single document. • Each User is considered a single document. 	

Data source connector	Supports crawling	Document definition	
		<ul style="list-style-type: none">• Each Task is considered a single document.• Each Idea is considered a single document.• Each Profile is considered a single document.• Each Solution is considered a single document.• Each Chatter is considered a single document.• Each Document (file) is considered a single document.• Each Custom entity (record) is considered a single document.• Each Knowledge article is considered a single document.	

Data source connector	Supports crawling	Document definition	
ServiceNow	<ul style="list-style-type: none"> • Incidents • Knowledge articles • Service catalog • Attachments 	<ul style="list-style-type: none"> • Each Incident is considered a single document. • Each Knowledge article is considered a single document. • Each Service catalog is considered a single document. • Each Attachment is considered a single document. 	
Slack	<ul style="list-style-type: none"> • Messages • Message attachments • Channel posts 	<ul style="list-style-type: none"> • Each Message is considered a single document. • Each Message attachment is considered a single document. • Each Channel post is considered a single document. 	

Data source connector	Supports crawling	Document definition	
Zendesk	<ul style="list-style-type: none"> • Tickets • Ticket comments • Ticket comment attachments • Articles • Article attachments • Article comments • Community topics • Community posts • Community post comments 	<ul style="list-style-type: none"> • Each Ticket is considered a single document. • Each Ticket comment is considered a single document. • Each Ticket comment attachment is considered a single document. • Each Article is considered a single document. • Each Article attachment is considered a single document. • Each Article comment is considered a single document. • Each Community topic is considered a single document. • Each Community post is considered a single document. • Each Community post comment is 	

Data source connector	Supports crawling	Document definition	
		considered a single document.	

Supported languages for Amazon Q Business

Amazon Q Business is optimized to respond in English. Amazon Q only indexes English language documents when you [connect a Amazon Q data source](#) or [directly upload documents](#) into your application. We recommend indexing only English language content.

Setting up for Amazon Q Business

Before you begin using Amazon Q Business for the first time, complete the following tasks.

Topics

- [Initial AWS account setup](#)
- [\(Optional\) Install the AWS CLI](#)
- [\(Optional\) Set up the AWS SDKs](#)
- [Consider AWS Regions and endpoints](#)
- [Set up required permissions](#)
- [Enable and configure an IAM Identity Center instance](#)

Initial AWS account setup

Sign up for an AWS account

If you do not have an AWS account, complete the following steps to create one.

To sign up for an AWS account

1. Open <https://portal.aws.amazon.com/billing/signup>.
2. Follow the online instructions.

Part of the sign-up procedure involves receiving a phone call and entering a verification code on the phone keypad.

When you sign up for an AWS account, an *AWS account root user* is created. The root user has access to all AWS services and resources in the account. As a security best practice, [assign administrative access to an administrative user](#), and use only the root user to perform [tasks that require root user access](#).

AWS sends you a confirmation email after the sign-up process is complete. At any time, you can view your current account activity and manage your account by going to <https://aws.amazon.com/> and choosing **My Account**.

Create an administrative user

After you sign up for an AWS account, secure your AWS account root user, enable AWS IAM Identity Center, and create an administrative user so that you don't use the root user for everyday tasks.

Secure your AWS account root user

1. Sign in to the [AWS Management Console](#) as the account owner by choosing **Root user** and entering your AWS account email address. On the next page, enter your password.

For help signing in by using root user, see [Signing in as the root user](#) in the *AWS Sign-In User Guide*.

2. Turn on multi-factor authentication (MFA) for your root user.

For instructions, see [Enable a virtual MFA device for your AWS account root user \(console\)](#) in the *IAM User Guide*.

Create an administrative user

1. Enable IAM Identity Center.

For instructions, see [Enabling AWS IAM Identity Center](#) in the *AWS IAM Identity Center User Guide*.

2. In IAM Identity Center, grant administrative access to an administrative user.

For a tutorial about using the IAM Identity Center directory as your identity source, see [Configure user access with the default IAM Identity Center directory](#) in the *AWS IAM Identity Center User Guide*.

Sign in as the administrative user

- To sign in with your IAM Identity Center user, use the sign-in URL that was sent to your email address when you created the IAM Identity Center user.

For help signing in using an IAM Identity Center user, see [Signing in to the AWS access portal](#) in the *AWS Sign-In User Guide*.

(Optional) Install the AWS CLI

The AWS Command Line Interface (AWS CLI) is a unified developer tool for managing AWS services, including Amazon Q.

- To install the AWS CLI, follow the instructions in [Installing the AWS Command Line Interface](#) in the *AWS Command Line Interface User Guide*.
- To configure the AWS CLI and set up a profile to call the AWS CLI, follow the instructions in [Configuring the AWS CLI](#) in the *AWS Command Line Interface User Guide*.
- To confirm that the AWS CLI profile is configured, run the following command:

```
aws configure --profile default
```

If your profile has been configured correctly, you will see output similar to the following:

```
AWS Access Key ID [*****52FQ]:  
AWS Secret Access Key [*****xgyZ]:  
Default region name [us-west-2]:  
Default output format [json]:
```

- To verify that the AWS CLI is configured for use with Amazon Q, run the following commands:

```
aws qbusiness help
```

If the AWS CLI is configured correctly, you will see a list of the supported AWS CLI commands for Amazon Q, Amazon Q runtime, and Amazon Q events.

(Optional) Set up the AWS SDKs

Download and install the AWS SDKs that you want to use. This guide provides examples for Python. For information about other AWS SDKs, see [Tools for Amazon Web Services](#).

The package for the Python SDK is called *Boto3*.

Before you run the following Python commands, you must first download and install [Python 3.6 or later](#) for your operating system. Support for Python 3.5 and earlier is deprecated.

If you don't have pip included in your Python Scripts directory, you can download [get-pip.py](#) and store this in your Scripts directory. You can also set your Python directory as a [Path or environment variable](#) using a terminal program.

To install Python, complete the following steps:

```
# Install the latest Boto3 release via pip
pip install boto3

# You can install a specific version of Boto3 for compatibility reasons
# Install Boto3 version 1.0 specifically
pip install boto3==1.0.0

# Make sure Boto3 is no older than version 1.15.0
pip install boto3>=1.15.0

# Avoid versions of Boto3 newer than version 1.15.3
pip install boto3<=1.15.3
```

To use Boto3, you must set up authentication credentials for your AWS account using the [IAM console](#).

Consider AWS Regions and endpoints

An *endpoint* is a URL that's the entry point for a web service. Each endpoint is associated with a specific AWS Region.

If you use a combination of the Amazon Q console, the AWS CLI, and the Amazon Q SDKs, pay attention to their default Regions. All Amazon Q components of a given application must be created in the same Region. Examples of a component include a retriever, an index, and a chat experience.

For the regions and endpoints supported by Amazon Q, see [Regions and Endpoints](#).

Set up required permissions

If you use Amazon Q through the AWS Management Console, required permissions are added on your behalf.

To use Amazon Q as an IAM user on the AWS CLI, or AWS SDK, you must attach the following permissions to allow Amazon Q to create and manage resources on your behalf:

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Action": "qbusiness:*",
    "Effect": "Allow",
    "Resource": "*"
  }]
}
```

For a complete list of IAM roles for Amazon Q, see [IAM roles for Amazon Q](#).

Enable and configure an IAM Identity Center instance

If you're planning to use IAM Identity Center as the user access manager for your Amazon Q application, we recommend enabling and pre-configuring an IAM Identity Center instance before you begin to create your Amazon Q application. If you do so, Amazon Q automatically detects—and connects to—your already configured IAM Identity Center instance.

If you're planning to use IAM Identity Center to connect your Amazon Q application to an Active Directory (AD) or external identity provider, creating a local IAM Identity Center instance and configuring it *before* you configure an Amazon Q application is recommended.

If you don't have an IAM Identity Center instance configured, and you want to use IAM Identity Center as your identity provider, you can also choose to create, connect, and minimally configure an IAM Identity Center instance for your Amazon Q application as part of the Amazon Q application creation process from the Amazon Q console.

You can add users to your IAM Identity Center instance from the Amazon Q console. When you add a new user to IAM Identity Center from the Amazon Q Business console, you need to make sure that the user is enabled for console access in your IAM Identity Center instance and their email ID is verified before they can log in to your Amazon Q web experience to chat. By default, a new user added to IAM Identity Center from the Amazon Q console isn't enabled. For more information on enabling users in IAM Identity Center, see [Creating an IAM user in your AWS account](#) in the IAM Identity Center User Guide.

You can't add groups to an IAM Identity Center instance from the Amazon Q console. If you want to add groups, Amazon Q will redirect you to the IAM Identity Center console to configure groups. To avoid this, you can configure groups in your IAM Identity Center instance before you create your Amazon Q app. Any groups already configured will be auto-detected by the Amazon Q console.

Your IAM Identity Center instance must be created in the same region as your Amazon Q application. To understand why this is important, see [Considerations for choosing an AWS Region](#) in the IAM Identity Center User Guide.

Important

Starting April 30, 2024, all new applications will need to use IAM Identity Center directly to manage user access. All existing Amazon Q applications will need to migrate to using IAM Identity Center for user management by July 31, 2024. We recommend you integrate any new application you're creating directly with IAM Identity Center.

IAM roles for Amazon Q Business

When you create an application or a web experience with Amazon Q Business, or connect a data source to it, Amazon Q needs access to the required AWS resources.

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create the Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role. The console displays roles that have the string **qbusiness** or **QBusiness** in the role name.

The following topics provide details for the required policies. If you create IAM roles using the Amazon Q console, these policies are created on your behalf.

Topics

- [IAM role for an Amazon Q Business application](#)
- [IAM role for an Amazon Q web experience](#)
- [IAM role for Amazon Q Business data source connectors](#)
- [IAM role for Amazon Q Business plugins](#)
- [IAM role for Lambda functions](#)
- [IAM role for an Amazon Kendra retriever](#)

IAM role for an Amazon Q Business application

When you create an Amazon Q Business application, you must provide Amazon Q with an IAM role with permissions to write to an Amazon CloudWatch log. You must also provide a trust policy that allows Amazon Q to assume the role. The following are the policies that must be provided.

To allow Amazon Q to access a CloudWatch log, use the following role policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AmazonQApplicationPutMetricDataPermission",
      "Effect": "Allow",
      "Action": [
        "cloudwatch:PutMetricData"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "cloudwatch:namespace": "AWS/QBusiness"
        }
      }
    },
    {
      "Sid": "AmazonQApplicationDescribeLogGroupsPermission",
      "Effect": "Allow",
      "Action": [
```

```

        "logs:DescribeLogGroups"
    ],
    "Resource": "*"
  },
  {
    "Sid": "AmazonQApplicationCreateLogGroupPermission",
    "Effect": "Allow",
    "Action": [
      "logs:CreateLogGroup"
    ],
    "Resource": [
      "arn:aws:logs:{{region}}:{{account_id}}:log-group:/aws/qbusiness/*"
    ]
  },
  {
    "Sid": "AmazonQApplicationLogStreamPermission",
    "Effect": "Allow",
    "Action": [
      "logs:DescribeLogStreams",
      "logs:CreateLogStream",
      "logs:PutLogEvents"
    ],
    "Resource": [
      "arn:aws:logs:{{region}}:{{account_id}}:log-group:/aws/qbusiness/*:log-
stream:*"
    ]
  }
]
}

```

To allow Amazon Q to assume a role, use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AmazonQApplicationPermission",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {

```

```

    "StringEquals": {
      "aws:SourceAccount": "{{account_id}}"
    },
    "ArnLike": {
      "aws:SourceArn": "arn:aws:qbusiness:{{region}}:{{account_id}}:application/*"
    }
  }
}
]
}

```

IAM role for an Amazon Q web experience

When you integrate your web experience with IAM Identity Center or deploy a web experience using an external identity provider, you must provide Amazon Q with an IAM role with permissions to write access relevant Amazon Q API operations. You must also provide a trust policy that allows Amazon Q to assume the role. The following are the policies that must be provided.

To allow Amazon Q to access the API operations required to integrate your application with IAM Identity Center or deploy your web experience using an external IdP, use the following role policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "QBusinessConversationPermission",
      "Effect": "Allow",
      "Action": [
        "qbusiness:Chat",
        "qbusiness:ChatSync",
        "qbusiness:ListMessages",
        "qbusiness:ListConversations",
        "qbusiness>DeleteConversation",
        "qbusiness:PutFeedback",
        "qbusiness:GetWebExperience",
        "qbusiness:GetApplication",
        "qbusiness:ListPlugins",
        "qbusiness:GetChatControlsConfiguration"
      ],
      "Resource": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
    }
  ]
}

```

```

    },
    {
      "Sid": "QBusinessKMSDecryptPermissions",
      "Effect": "Allow",
      "Action": [
        "kms:Decrypt"
      ],
      "Resource": [
        "arn:aws:kms:{{region}}:{{account_id}}:key/[{{key_id}}]"
      ],
      "Condition": {
        "StringLike": {
          "kms:ViaService": [
            "qbusiness.{{region}}.amazonaws.com"
          ]
        }
      }
    }
  ]
}

```

To allow Amazon Q to assume a role, use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "QBusinessTrustPolicy",
      "Effect": "Allow",
      "Principal": {
        "Service": "application.qbusiness.amazonaws.com"
      },
      "Action": [
        "sts:AssumeRole",
        "sts:SetContext"
      ],
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "{{source_account}}"
        },
        "ArnEquals": {
          "aws:SourceArn": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/{{application_id}}"
        }
      }
    }
  ]
}

```



```
}
  }
}
]
```

IAM role for Amazon Q Business data source connectors

You can use either the Amazon Q Business console or the [CreateDataSource](#) API operation to connect your data source. However, you must first provide Amazon Q with an IAM role that has permissions to access the data source resources.

If you use the console, you can either create an IAM role when you connect your data source to Amazon Q or use an existing role. If you use the `CreateDataSource` API operation, you must provide the Amazon Resource Name (ARN) of an existing IAM role.

The specific permissions required depend on the data source. At a minimum, your IAM role must include the following:

- Permission to access the [BatchPutDocument](#) and [BatchDeleteDocument](#) API operations in order to ingest documents.
- Permission to access the User Store APIs needed to ingest access control and identity information from documents.

To allow Amazon Q to connect to your data source, use the following least-permissions role policy:

Note

This policy assumes your data source doesn't use any authentication.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AllowsAmazonQToIngestDocuments",
    "Effect": "Allow",
    "Action": [
      "qbusiness:BatchPutDocument",
```

```

        "qbusiness:BatchDeleteDocument"
    ],
    "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
  },
  {
    "Sid": "AllowsAmazonQToIngestPrincipalMapping",
    "Effect": "Allow",
    "Action": [
      "qbusiness:PutGroup",
      "qbusiness:CreateUser",
      "qbusiness>DeleteGroup",
      "qbusiness:UpdateUser",
      "qbusiness:ListGroups"
    ],
    "Resource": [
      "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}",
      "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}",
      "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}/data-source/*"
    ]
  }
]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToAssumeRoleForServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "{{source_account}}"
        }
      },
    }
  ]
}

```

```

    "ArnLike": {
      "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
    }
  }
}
]
}

```

If your data source uses authentication, you must add the following policy to your IAM role to allow Amazon Q to access your AWS Secrets Manager secret:

```

{
  "Sid": "AllowsAmazonQToGetSecret",
  "Effect": "Allow",
  "Action": [
    "secretsmanager:GetSecretValue"
  ],
  "Resource": [
    "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
  ]
}

```

If you are using an Amazon VPC, you must add the following VPC access permissions to your policy:

```

{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AllowsAmazonQToCreateAndDeleteNI",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2:DeleteNetworkInterface"
    ],
    "Resource": [
      "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[subnet_ids]",
      "arn:aws:ec2:{{region}}:{{account_id}}:security-group/
[[security_group]]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",

```

```

    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2>DeleteNetworkInterface"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:RequestTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
      },
      "ForAllValues:StringEquals": {
        "aws:TagKeys": [
          "AMAZON_Q"
        ]
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateTags",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateTags"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringEquals": {
        "ec2:CreateAction": "CreateNetworkInterface"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterfacePermission"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:ResourceTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
      }
    }
  }
}

```

```

    },
    {
      "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
      "Effect": "Allow",
      "Action": [
        "ec2:DescribeNetworkInterfaces",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeNetworkInterfaceAttribute",
        "ec2:DescribeVpcs",
        "ec2:DescribeRegions",
        "ec2:DescribeNetworkInterfacePermissions",
        "ec2:DescribeSubnets"
      ],
      "Resource": "*"
    }
  ]
}

```

If your Secrets Manager secret is encrypted, you must add permissions for AWS KMS key to decrypt the username and password secret stored by Secrets Manager:

```

{
  "Effect": "Allow",
  "Action": [
    "kms:Decrypt"
  ],
  "Resource": [
    "arn:aws:kms:your-region:your-account-id:key/key-id"
  ],
  "Condition": {
    "StringLike": {
      "kms:ViaService": [
        "secretsmanager.*.amazonaws.com"
      ]
    }
  }
}

```

If your Amazon Q data source connector needs access to an object stored in an Amazon S3 bucket (such as an SSL certificate), you must add the following permissions to your IAM role:

Note

Check that the file path to the object in your Amazon S3 bucket is of the following format:
s3://BucketName/FolderName/FileName.extension.

```
{
  "Sid": "AllowsAmazonQToGetS3Objects",
  "Action": [
    "s3:GetObject"
  ],
  "Resource": [
    "arn:aws:s3:::{{input_bucket_name}}/*"
  ],
  "Effect": "Allow",
  "Condition": {
    "StringEquals": {
      "aws:ResourceAccount": "{{account_id}}"
    }
  }
}
```

IAM role for Amazon S3 data sources

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q Business resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

When you use an Amazon S3 bucket as a data source, you must provide a role that has permissions to:

- Access your Amazon S3 bucket.
- Permission to access the [BatchPutDocument](#) and [BatchDeleteDocument](#) API operations in order to ingest documents.
- Permission to access the Principal Store APIs needed to ingest access control and identity information from documents.

To allow Amazon Q to use an Amazon S3 bucket as a data source, use the following role policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToGetObjectfromS3",
      "Action": [
        "s3:GetObject"
      ],
      "Resource": [
        "arn:aws:s3:::{{input_bucket_name}}/*"
      ],
      "Effect": "Allow",
      "Condition": {
        "StringEquals": {
          "aws:ResourceAccount": "{{account_id}}"
        }
      }
    },
    {
      "Sid": "AllowsAmazonQToListS3Buckets",
      "Action": [
        "s3:ListBucket"
      ],
      "Resource": [
        "arn:aws:s3:::{{input_bucket_name}}"
      ],
      "Effect": "Allow",
      "Condition": {
        "StringEquals": {
          "aws:ResourceAccount": "{{account_id}}"
        }
      }
    },
    {
      "Sid": "AllowsAmazonQToIngestDocuments",
      "Effect": "Allow",
      "Action": [
        "qbusiness:BatchPutDocument",
        "qbusiness:BatchDeleteDocument"
      ],
      "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
        {{application_id}}/index/{{index_id}}"
    }
  ]
}
```

```

    },
    {
      "Sid": "AllowsAmazonQToCallPrincipalMappingAPIs",
      "Effect": "Allow",
      "Action": [
        "qbusiness:PutGroup",
        "qbusiness:CreateUser",
        "qbusiness>DeleteGroup",
        "qbusiness:UpdateUser",
        "qbusiness:ListGroups"
      ],
      "Resource": [
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}",
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}",
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}/data-source/*"
      ]
    }
  ]
}

```

If the documents in the Amazon S3 bucket are encrypted, you must provide the following permissions to use the AWS KMS key to decrypt the documents:

```

{
  "Sid": "AllowsAmazonQToDecryptSecret",
  "Effect": "Allow",
  "Action": [
    "kms:Decrypt"
  ],
  "Resource": [
    "arn:aws:kms:{{region}}:{{account_id}}:key/[key_id]"
  ],
  "Condition": {
    "StringLike": {
      "kms:ViaService": [
        "secretsmanager.*.amazonaws.com"
      ]
    }
  }
}

```


If you are using an Amazon VPC, you must add the following VPC access permissions to your policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToGetObjectfromS3",
      "Action": [
        "s3:GetObject"
      ],
      "Resource": [
        "arn:aws:s3:::{{input_bucket_name}}/*"
      ],
      "Effect": "Allow",
      "Condition": {
        "StringEquals": {
          "aws:ResourceAccount": "{{account_id}}"
        }
      }
    },
    {
      "Sid": "AllowsAmazonQToListS3Buckets",
      "Action": [
        "s3:ListBucket"
      ],
      "Resource": [
        "arn:aws:s3:::{{input_bucket_name}}"
      ],
      "Effect": "Allow",
      "Condition": {
        "StringEquals": {
          "aws:ResourceAccount": "{{account_id}}"
        }
      }
    },
    {
      "Sid": "AllowsAmazonQToIngestDocuments",
      "Effect": "Allow",
      "Action": [
        "qbusiness:BatchPutDocument",
        "qbusiness:BatchDeleteDocument"
      ],
    }
  ]
}
```

```

    "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
    {{application_id}}/index/{{index_id}}"
  },
  {
    "Sid": "AllowsAmazonQToCallPrincipalMappingAPIs",
    "Effect": "Allow",
    "Action": [
      "qbusiness:PutGroup",
      "qbusiness:CreateUser",
      "qbusiness>DeleteGroup",
      "qbusiness:UpdateUser",
      "qbusiness>ListGroups"
    ],
    "Resource": [
      "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}",
      "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
      index/{{index_id}}",
      "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
      index/{{index_id}}/data-source/*"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateAndDeleteENI",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2>DeleteNetworkInterface"
    ],
    "Resource": [
      "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
      "arn:aws:ec2:{{region}}:{{account_id}}:security-group/[{{security_group}}]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateDeleteENI",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2>DeleteNetworkInterface"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:RequestTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
      }
    }
  }
}

```

```

    },
    "ForAllValues:StringEquals": {
      "aws:TagKeys": [
        "AMAZON_Q"
      ]
    }
  }
},
{
  "Sid": "AllowsAmazonQToCreateTags",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateTags"
  ],
  "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
  "Condition": {
    "StringEquals": {
      "ec2:CreateAction": "CreateNetworkInterface"
    }
  }
},
{
  "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateNetworkInterfacePermission"
  ],
  "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
  "Condition": {
    "StringLike": {
      "aws:ResourceTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
    }
  }
},
{
  "Sid": "AllowsAmazonQToConnectToVPC",
  "Effect": "Allow",
  "Action": [
    "ec2:DescribeNetworkInterfaces",
    "ec2:DescribeAvailabilityZones",
    "ec2:DescribeNetworkInterfaceAttribute",
    "ec2:DescribeVpcs",
    "ec2:DescribeRegions",
    "ec2:DescribeNetworkInterfacePermissions",

```

```

    "ec2:DescribeSubnets"
  ],
  "Resource": "*"
}
]
}

```

To allow Amazon Q to assume a role, use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToAssumeRoleForServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "{{source_account}}"
        },
        "ArnLike": {
          "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
        }
      }
    }
  ]
}

```

IAM role for Amazon Q Business plugins

To successfully connect Amazon Q Business to a plugin, you need to give Amazon Q Business the following permissions using a service access role:

- Permission to access your Secrets Manager secret to get the credentials you use to log in to the third party service instance you are creating a plugin for.
- **(Optional)** Permission to access the customer managed AWS KMS key used to encrypt the content of your Secrets Manager secret.

Amazon Q assumes this role to access your third party service instance credentials.

If you use the console and choose to create a new IAM role, Amazon Q creates the IAM role for you. If you use the console and choose to use an existing secret, or you use the API, make sure your secret contains the following permissions.

The following is the service access IAM role required:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowQBusinessToGetSecretValue",
      "Effect": "Allow",
      "Action": [
        "secretsmanager:GetSecretValue"
      ],
      "Resource": [
        "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[[secret_id]]"
      ]
    }
  ]
}
```

To allow Amazon Q to assume a role, use the following trust policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "QBusinessApplicationTrustPolicy",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "{{source_account}}"
        },
        "ArnLike": {
          "aws:SourceArn": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}"
        }
      }
    }
  ]
}
```

```
    }  
  }  
}  
]  
}
```

IAM role for Lambda functions

Custom document enrichment (CDE) is an Amazon Q Business feature that you can use to manipulate your document content and document attributes. When you use the Lambda functions for CDE, you need an IAM role for the following:

- A role for `PreExtractionHookConfiguration` with permissions to run `PreExtractionHookConfiguration` and to access the Amazon S3 bucket when you use `PreExtractionHookConfiguration`.
- A role for `PostExtractionHookConfiguration` with permissions to run `PostExtractionHookConfiguration` and to access the Amazon S3 bucket when you use `PostExtractionHookConfiguration`.

Important

IAM roles for Custom Document Enrichment (CDE) Lambda functions should belong to the same account as the account using [BatchPutDocument](#) API operation or the [CreateDataSource](#) operation to configure CDE.

Both AWS Identity and Access Management (IAM) roles must have the permissions to:

- Run `PreExtractionHookConfiguration` and/or `PostExtractionHookConfiguration`. To apply advanced alterations of your document metadata and content during the ingestion process, configure a Lambda function for `PreExtractionHookConfiguration` and/or `PostExtractionHookConfiguration`.
- (Optional) If you choose to activate Server Side Encryption for your Amazon S3 bucket, you must provide permissions to use the AWS KMS key customer to encrypt and decrypt the objects stored in your Amazon S3 bucket.

A role policy to allow Amazon Q to run PreExtractionHookConfiguration with encryption for your Amazon S3 bucket.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Action": [
      "s3:GetObject",
      "s3:PutObject"
    ],
    "Resource": [
      "arn:aws:s3:::bucket-name",
      "arn:aws:s3:::bucket-name/*"
    ],
    "Effect": "Allow"
  },
  {
    "Action": [
      "s3:ListBucket"
    ],
    "Resource": [
      "arn:aws:s3:::bucket-name"
    ],
    "Effect": "Allow"
  },
  {
    "Effect": "Allow",
    "Action": [
      "kms:Decrypt",
      "kms:GenerateDataKey"
    ],
    "Resource": [
      "arn:aws:kms:your-region:your-account-id:key/key-id"
    ]
  },
  {
    "Effect": "Allow",
    "Action": [
      "lambda:InvokeFunction"
    ],
    "Resource": "arn:aws:lambda:your-region:your-account-id:function:pre-extraction-lambda-function"
  }
}
```

```
    ]
  }
}
```

An role policy to allow Amazon Q to run PreExtractionHookConfiguration without encryption.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Action": [
      "s3:GetObject",
      "s3:PutObject"
    ],
    "Resource": [
      "arn:aws:s3:::bucket-name",
      "arn:aws:s3:::bucket-name/*"
    ],
    "Effect": "Allow"
  },
  {
    "Action": [
      "s3:ListBucket"
    ],
    "Resource": [
      "arn:aws:s3:::bucket-name"
    ],
    "Effect": "Allow"
  },
  {
    "Effect": "Allow",
    "Action": [
      "lambda:InvokeFunction"
    ],
    "Resource": "arn:aws:lambda:your-region:your-account-id:function:pre-
extraction-lambda-function"
  }
  ]
}
```

A role policy to allow Amazon Q to run PostExtractionHookConfiguration with encryption for your Amazon S3 bucket.


```

{
  "Version": "2012-10-17",
  "Statement": [{
    "Action": [
      "s3:GetObject",
      "s3:PutObject"
    ],
    "Resource": [
      "arn:aws:s3:::bucket-name",
      "arn:aws:s3:::bucket-name/*"
    ],
    "Effect": "Allow"
  },
  {
    "Action": [
      "s3:ListBucket"
    ],
    "Resource": [
      "arn:aws:s3:::bucket-name"
    ],
    "Effect": "Allow"
  },
  {
    "Effect": "Allow",
    "Action": [
      "kms:Decrypt",
      "kms:GenerateDataKey"
    ],
    "Resource": [
      "arn:aws:kms:your-region:your-account-id:key/key-id"
    ]
  },
  {
    "Effect": "Allow",
    "Action": [
      "lambda:InvokeFunction"
    ],
    "Resource": "arn:aws:lambda:your-region:your-account-id:function:post-
extraction-lambda-function"
  }
  ]
}

```

An role policy to allow Amazon Q to run PostExtractionHookConfiguration without encryption.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Action": [
      "s3:GetObject",
      "s3:PutObject"
    ],
    "Resource": [
      "arn:aws:s3:::bucket-name",
      "arn:aws:s3:::bucket-name/*"
    ],
    "Effect": "Allow"
  },
  {
    "Action": [
      "s3:ListBucket"
    ],
    "Resource": [
      "arn:aws:s3:::bucket-name"
    ],
    "Effect": "Allow"
  },
  {
    "Effect": "Allow",
    "Action": [
      "lambda:InvokeFunction"
    ],
    "Resource": "arn:aws:lambda:your-region:your-account-id:function:post-extraction-lambda-function"
  }
]
```

We recommend that you include `aws:sourceAccount` and `aws:sourceArn` in the trust policy. Their inclusion limits permissions and securely checks if `aws:sourceAccount` and `aws:sourceArn` are the same values as provided in the IAM role policy for the `sts:AssumeRole` action. This approach prevents unauthorized entities from accessing your IAM roles and their permissions. For more information, see [confused deputy problem](#) in the *IAM User Guide*.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": [
          "qbusiness.amazonaws.com"
        ]
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "your-account-id"
        },
        "StringLike": {
          "aws:SourceArn": "arn:aws:qbusiness:your-region:your-account-id:application/
<application-id>/index/<index-id>"
        }
      }
    }
  ]
}
```

IAM role for an Amazon Kendra retriever

When you use an Amazon Kendra index as a retriever, you must provide Amazon Q Business with an IAM role with permissions to access Amazon Kendra. You must also provide a trust policy that allows Amazon Q to assume the role. The following are the policies that must be provided.

To allow Amazon Q to access a CloudWatch log, use the following policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "KendraRetrieveAccess",
      "Effect": "Allow",
      "Action": [
        "kendra:Retrieve",
        "kendra:DescribeIndex"
      ]
    }
  ]
}
```

```

        ],
        "Resource": "arn:aws:kendra:{{region}}:{{source_account}}:index/
{{indexId}}"
    }
]
}

```

To allow Amazon Q to assume a role, use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AmazonQKendraAccessPermission",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "{{source_account}}"
        },
        "ArnEquals": {
          "aws:SourceArn": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{applicationId}}"
        }
      }
    }
  ]
}

```

Creating a sample Amazon Q Business application

You can create a sample Amazon Q Business application to test the capabilities of Amazon Q before creating a [fully-configured application](#). A sample application supports only [upload file and chat](#) conversations, is powered by an Amazon Q [native retriever](#), and doesn't have to be connected to [Amazon Q data sources](#).

Getting started with your sample application is a two-step process. First, you create an Amazon Q Business application and integrate it with IAM Identity Center from within the Amazon Q console. Then, you add either existing or new users and groups to your Amazon Q Business application. When your application is successfully created, Amazon Q generates a web experience login URL. Any user you've added to your application and enabled can log in and chat with Amazon Q.

You can't connect data sources or choose a retriever for a sample Amazon Q Business application when you create it. However, you can choose to select a retriever and connect data sources when you [update it](#).

Your IAM Identity Center instance must be created in the same region as your Amazon Q application. To understand why this is important, see [Considerations for choosing an AWS Region](#) in the IAM Identity Center User Guide.

Topics

- [Step 1: Create a sample application](#)
- [Step 2: Add users and groups](#)
- [Managing a sample application](#)

Step 1: Create a sample application

This section guides you through the process of creating a sample Amazon Q application. To do this, you can use the Amazon Q console, the AWS Command Line Interface (AWS CLI), and the Amazon Q API operations.

Console

To create an sample application

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.

2. From the **How it works** menu, from **Experiment with a sample – optional**, choose **Try quick application**.
3. On the **Create application** page, for **Application settings**, enter the following information for your Amazon Q application:

- **Application name** – A name for your Amazon Q Business application for easy identification. This name is only visible in the AWS Management Console. The name can include hyphens (-), but not spaces, and can have a maximum of 1,000 alphanumeric characters.
- **Service access** – An IAM role for Amazon Q Business to allow it to access the AWS resources it needs to create your application. You can choose to use an existing role or create a new role.

 **Note**

For more information about example service roles, see [IAM role for an Amazon Q Business application](#).

- **Service role name** – A name for the service (IAM) role you created for easy identification on the console.
- **Encryption** – Amazon Q encrypts your data by default using AWS managed AWS KMS keys. To customize your encryption settings, select **Customize encryption settings (advanced)**. Then, you can choose to use an existing AWS KMS key or create a new one. To learn more, see [Data encryption](#).

 **Important**

If you choose to use a customer managed key, you must provision at least 10 index storage units when you [create an Amazon Q retriever](#).

4. In **Connect Amazon Q to IAM Identity Center**, you will see the following options based on whether you have an IAM Identity Center instance already configured, or need to create one.
 1. If you don't have an IAM Identity Center instance configured, you see the following:
 - The region your Amazon Q application is in. This is so you can make sure that the region for your Amazon Q application and IAM Identity Center instance match.

- **Specify tags for IAM Identity Center** – Add tags to keep track of your IAM Identity Center instance.
- **Create IAM Identity Center** – Select to create a minimally-configured IAM Identity Center instance. The console will display an ARN for your newly created resource after it's created.

 **Note**

You can't add groups to your application from your Amazon Q console unless you already have an IAM Identity Center instance with groups configured. When you add a new user to IAM Identity Center from the Amazon Q Business console, you need to make sure that the user is enabled in your IAM Identity Center instance and their email ID is verified before they can log in to your Amazon Q Business web experience to chat.

2. If you have *both* an IAM Identity Center organization instance and an account instance configured, your instances will be auto-detected, and you see the following options:
 - **[Connect to organization instance of IAM Identity Center](#)** – Select this option to manage access to Amazon Q by assigning users and groups from the Identity Center directory for your organization.
 - **[Connect to account instance of IAM Identity Center](#)** – Select this option to manage access to Amazon Q by assigning existing users and groups from your Identity Center directory.
 - The region your Amazon Q application is in. This is so you can make sure that the region for your Amazon Q application and IAM Identity Center instance match.
 - **IAM Identity Center** – The ARN for your IAM Identity Center instance.
3. If you have an IAM Identity Center account instance configured, your account instance will be auto-detected and you will see the following:
 - The region your Amazon Q application is in. This is so you can make sure that the region for your Amazon Q application and IAM Identity Center instance match.
 - **IAM Identity Center** – The ARN for your IAM Identity Center instance.
4. If you have an IAM Identity Center organization instance configured, you will see a message asking you to tell your admin to give you access to IAM Identity Center. You will need access to IAM Identity Center before you can proceed.

5. **Tags – optional** – To add tags to your Amazon Q application and web experience, select **Add new tag**. Then, enter the following information for each tag:

- **Key** – Add a key for your tag.
- **Value - optional** – An optional value for your tag.

For more information about using tags with Amazon Q, see [Tags](#).

6. To start creating your application, choose **Create**.

AWS CLI

To configure an Amazon Q application

```
aws qbusiness create-application \  
--display-name application-name \  
--identity-center-instance-arn identity-center-instance-arn \  
--role-arn roleArn \  
--description application-description \  
--encryption-configuration kmsKeyId=<kms-key-id> \  
--attachments-configuration attachmentsControlMode=ENABLED
```

Step 2: Add users and groups

An Amazon Q Business sample application requires you to use IAM Identity Center to manage user access. While it's recommended to have an IAM Identity Center instance configured (with users and groups added) before you start, you can also choose to create and configure an IAM Identity Center instance for your sample Amazon Q application from within the Amazon Q console.

You can add users to your IAM Identity Center instance from the Amazon Q Business console. When you add a new user to IAM Identity Center from the Amazon Q Business console, you need to make sure that the user is enabled in your IAM Identity Center instance and their email ID is verified before they can log in to your Amazon Q Business web experience to chat.

New users added to IAM Identity Center from the Amazon Q console must accept the invitation to IAM Identity Center they receive in their email and set their passwords before they can access your Amazon Q application through their web experience URL. By default, a new user added to IAM

Identity Center from the Amazon Q console isn't enabled. For more information on enabling users in IAM Identity Center, see [Add users](#) in the IAM Identity Center User Guide.

You can't add groups to your application from your Amazon Q console unless you already have an IAM Identity Center instance with groups configured. You must add and assign at least one user for your Amazon Q application.

On successful completion, Amazon Q Business returns a web experience URL that you can share with the end users you added to your application. For an overview of your application's capabilities, see [Using a web experience](#).

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To add users and groups to your Amazon Q application

1. To add users, in **Add groups and users**, choose **Users**, and then choose **Add users**. Then, depending on how you're integrating Amazon Q with IAM Identity Center, do the following:
 - a. If you're using a pre-configured IAM Identity Center instance with users and groups already added, Amazon Q detects the users you have configured in IAM Identity Center. You can choose to assign users from your IAM Identity Center directory.
 - i. In this case, in the **Assign users and groups** dialog box that opens, type and select the name of the users you want to assign.
 - ii. Then, select the user you want to add and then select **Assign**.
 - b. If you've created a minimally-configured IAM Identity Center instance from within the Amazon Q console for your Amazon Q application, you can add users to your IAM Identity Center instance and then assign new users to your Amazon Q application in this step.
 - i. In this case, in the **Add or assign users and groups** dialog box that opens, select **Add new users** and select **Next**.
 - ii. In **Add new users**, select **Add new users**.
 - iii. In **Add new users**, enter the following information:

- **Username** – A username is required for an user to sign into the AWS access portal. You can't change the username later. Maximum length 128 characters. Can only contain alphanumeric characters or any of the following: +=,.,@-_
 - **First name** – First name of user.
 - **Last name** – Last name of user.
 - **Email address** – Email address of user.
 - **Confirm email address** – Enter email address again to confirm it.
 - **Display name** – The display name assigned to your user.
- iv. Then, select **Next**. The **Add new users** dialog box displays the user who you've just added.
 - v. Select **Add**.

The new user you add will receive a notification to the email you entered for them asking them to accept your invitation to IAM Identity Center. You will also have to go to the IAM Identity Center console and send them an email verification request. Your user will have to complete the email verification before they can successfully log in to the web experience URL for your Amazon Q application. For more information on enabling users in IAM Identity Center, see [Add users](#) in the IAM Identity Center User Guide.

- vi. From the **Add or assign users and groups** dialog box that opens, choose **Assign existing users and groups** and enter and select the name of the user you added in the last step. Then, choose **Assign**.

When the assignment process successfully completes, the Amazon Q console will display the name, username, and user ID of the user assigned to your Amazon Q application.

2. To add groups, in **Add groups and users**, choose **Groups**, and then choose **Add groups**. Then, depending on how you're integrating Amazon Q with IAM Identity Center, do the following:
 - a. If you're using a pre-configured IAM Identity Center instance with users and groups already added, Amazon Q detects the groups you have configured in IAM Identity Center. You can choose to assign groups from your IAM Identity Center directory.

- i. In this case, in the **Assign users and groups** dialog box that opens, type and select the name of the groups you want to assign.
- ii. Then, select the group you want to add and select **Assign**.

 **Note**

As an IAM Identity Center user, you also have the option to add and adding a new user to your app from within the Amazon Q console. See the following step for details on how to do so. However, you can't create groups from within the Amazon Q console.

- b. If you've created a minimally-configured IAM Identity Center instance from within the Amazon Q console for your Amazon Q application, you can't add groups to your IAM Identity Center instance from the Amazon Q console.
 - i. In this case, in the **Add or assign users and groups** dialog box that opens, select **Manage users and groups from IAM Identity Center**.
 - ii. From the left navigation menu, choose **Groups**.
 - iii. From **Groups**, choose **Create group**. Then, enter the following information:
 - **Group name** – A name for your group.
 - **Description** – An optional group description.
 - iv. In **Add users to groups**, add users you want to the group.

 **Note**

The users must already exist in your IAM Identity Center instance. If they don't, you can select **Add new user** and follow the steps to add a new user to the group you're creating.

- v. Then, select **Create group**. The console will display a group summary page.
3. In **Web experience service access**, enter the following information:
 - For **Choose a method to authorize Amazon Q** – A service access role assumed by end users when they sign in to your web experience that grants them permission to start and manage conversations Amazon Q. You can choose to use an existing role or create a new role.

- **Service role name** – A name for the service role you created for easy identification on the console.
4. Select **Save**.
 5. Select **Create application**.

AWS CLI

To add users to an application

```
aws sso-admin create-application-assignment \  
--application-arn idc-app-arn \  
--principal-id idc-user-ID \  
--principal-type USER
```

To add groups to an application

```
aws sso-admin create-application-assignment \  
--application-arn idc-app-arn \  
--principal-id idc-group-ID \  
--principal-type GROUP
```

Managing a sample application

You can manage users and groups for your IAM Identity Center-integrated application from your Amazon Q application settings page. To access application settings, login to the Amazon Q console, choose your application from the **Applications** page, and then scroll down to **Users and groups**.

You can add new users to your application from the Amazon Q console, or choose to add users already existing in your IAM Identity Center directory.

You can only add groups already configured in your IAM Identity Center directory to your Amazon Q application using the Amazon Q console.

You can delete both users and groups from your Amazon Q application using the Amazon Q console. Deleting users and groups from your Amazon Q application using the Amazon Q console doesn't delete them from your IAM Identity Center directory.

To manage users and groups programmatically for your Amazon Q application, refer to the [IAM Identity Center CLI Reference](#) and the [Identity Store API Reference](#).

To learn more about managing your sample application, see [Managing Amazon Q Business applications](#).

Configuring an Amazon Q Business application

As the first step towards creating an Amazon Q Business chat application for your end users, you configure an Amazon Q application. Then, you can optionally enhance it by customizing the end user experience. After this, you select and create a retriever, and connect and configure the data sources.

This section guides you through the process of creating and configuring an Amazon Q application. To create an application, you can use the Amazon Q console, the AWS Command Line Interface (AWS CLI), and the Amazon Q API operations.

As a prerequisite, make sure that you complete the [setting up](#) tasks. If you're using the AWS CLI or the API, make sure that you created the required [IAM roles](#).

After you finish creating your application, you can customize and preview the web experience that it will power.

Important

Starting April 30, 2024, all new applications will need to use IAM Identity Center directly to manage user access. All existing Amazon Q applications using [legacy identity management](#) will need to migrate to using IAM Identity Center for user management by July 31, 2024. We recommend you integrate any new application you're creating directly with IAM Identity Center.

Note

During Preview, an Amazon Q application supports only 50 end users. If you need more capacity, contact [Support](#).

Topics

- [Creating an Amazon Q Business application](#)
- [Creating and selecting a retriever for an Amazon Q Business application](#)
- [Connecting data sources to an Amazon Q Business application](#)
- [Adding users and groups to an Amazon Q application](#)

Creating an Amazon Q Business application

To create an Amazon Q Business application, you can use either the AWS Management Console or the Amazon Q API.

Before you begin to create an Amazon Q application, make sure that you complete the [setting up](#) tasks. If you're using the AWS CLI or the Amazon Q API, make sure that you created the required [IAM roles](#).

After you create an application, you can create your Amazon Q web experience. How you create the web experience depends on whether you use the AWS Management Console or the Amazon Q APIs.

- **AWS Management Console** – If you use the console to create an application, the web experience is created automatically.
- **Amazon Q API** – If you use the [CreateApplication](#) API operation to create an application, use the [CreateWebExperience](#) API operation to create your web experience.

The following tabs provide a procedure for using IAM Identity Center as the user access manager for your Amazon Q application using the AWS Management Console and code examples for using the AWS CLI.

Important

Your Amazon Q application must be created in the same region as your IAM Identity Center instance.

Console

To configure an Amazon Q application

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. For **Create Amazon Q application**, choose **Get started**.
3. For **Applications**, choose **Create application**. The console will display a **Select access management method for application** dialog box.

4. In **Select access management method for application**, choose **IAM Identity Center (Recommended)**, and then select **Ok**. Choosing this option allows you to use IAM Identity Center as your AWS gateway to the identity provider of your choice.
5. For **Application settings**, enter the following information for your Amazon Q application:
 - **Application name** – A name for your Amazon Q Business application for easy identification. This name is only visible in the AWS Management Console. The name can include hyphens (-), but not spaces, and can have a maximum of 1,000 alphanumeric characters.
 - **Service access** – An IAM role for Amazon Q Business to allow it to access the AWS resources it needs to create your application. You can choose to use an existing role or create a new role.

 **Note**

For more information about example service roles, see [IAM role for an Amazon Q Business application](#).

- **Service role name** – A name for the service (IAM) role you created for easy identification on the console.
- **Encryption** – Amazon Q encrypts your data by default using AWS managed AWS KMS keys. To customize your encryption settings, select **Customize encryption settings (advanced)**. Then, you can choose to use an existing AWS KMS key or create a new one. To learn more, see [Data encryption](#).

 **Important**

If you choose to use a customer managed key, you must provision at least 10 index storage units when you [create an Amazon Q retriever](#).

6. In **Connect Amazon Q to IAM Identity Center**, you will see the following options based on whether you have an IAM Identity Center instance already configured, or need to create one.
 1. If you don't have an IAM Identity Center instance configured, you see the following:
 - The region your Amazon Q application is in. This is so you can make sure that the region for your Amazon Q application and IAM Identity Center instance match.

- **Specify tags for IAM Identity Center** – Add tags to keep track of your IAM Identity Center instance.
- **Create IAM Identity Center** – Select to create a minimally-configured IAM Identity Center instance. The console will display an ARN for your newly created resource after it's created.

 **Note**

If you plan to connect your IAM Identity Center to an Active Directory or external identity provider we recommend cancelling this setup and configuring IAM Identity Center from the IAM Identity Center console. If you're managing users and groups in one identity source, changing to a different identity source might remove all user and group assignments.

If you plan to add groups to your application using your minimally-configured IAM Identity Center instance, we recommend configuring these groups in IAM Identity Center before you create your application. If you don't have already configured IAM Identity Center groups, Amazon Q will redirect you to the IAM Identity Center console to configure groups before you can add them to your application.

2. If you have *both* an IAM Identity Center organization instance and an account instance configured, your instances will be auto-detected, and you see the following options:
 - **[Connect to organization instance of IAM Identity Center](#)** – Select this option to manage access to Amazon Q by assigning users and groups from the Identity Center directory for your organization.
 - **[Connect to account instance of IAM Identity Center](#)** – Select this option to manage access to Amazon Q by assigning existing users and groups from your Identity Center directory.
 - The region your Amazon Q application is in. This is so you can make sure that the region for your Amazon Q application and IAM Identity Center instance match.
 - **IAM Identity Center** – The ARN for your IAM Identity Center instance.
3. If you have an IAM Identity Center account instance configured, your account instance will be auto-detected and you will see the following:
 - The region your Amazon Q application is in. This is so you can make sure that the region for your Amazon Q application and IAM Identity Center instance match.

- **IAM Identity Center** – The ARN for your IAM Identity Center instance.
4. If you have an IAM Identity Center organization instance configured, you will see a message asking you to tell your admin to give you access to IAM Identity Center. You will need access to IAM Identity Center before you can proceed.
 7. **Tags – optional** – To add tags to your Amazon Q application and web experience, select **Add new tag**. Then, enter the following information for each tag:
 - **Key** – Add a key for your tag.
 - **Value - optional** – An optional value for your tag.

For more information about using tags with Amazon Q, see [Tags](#).

8. To start creating your application, choose **Create**.

AWS CLI

To configure an Amazon Q application

```
aws qbusiness create-application \  
--display-name application-name \  
--identity-center-instance-arn identity-center-instance-arn \  
--role-arn roleArn \  
--description application-description \  
--encryption-configuration kmsKeyId=<kms-key-id> \  
--attachments-configuration attachmentsControlMode=ENABLED
```

Managing Amazon Q Business applications

To manage an Amazon Q Business application, you can take the following actions:

Actions

- [Deleting an application](#)
- [Getting application properties](#)
- [Listing applications](#)
- [Updating an application](#)

Deleting an application

To delete an Amazon Q Business application, you can use the console or the [DeleteApplication](#) API operation.

The following tabs provide a procedure for the console and code examples for the AWS CLI.

Console

To delete an Amazon Q application

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. For **Applications**, choose **Actions**.
3. Choose **Delete**.
4. In the dialog box that opens, type **Delete** to confirm deletion, and then choose **Delete**.

You are returned to the service console while your application is deleted. When the deletion process is complete, the console displays a message confirming successful deletion.

AWS CLI

To delete an Amazon Q application

```
aws qbusiness delete-application \  
--application-id application-id
```

Getting application properties

To get the properties of an Amazon Q Business application, you can use the console or the [GetApplication](#) API operation.

The following tabs provide a procedure for the console and code examples for the AWS CLI.

Console

To get properties of an Amazon Q application

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. For **Applications**, select the name of your application from the list of applications.
3. On **Application settings**, the following properties are available:
 - **Application name** – The name that you chose for your application.
 - **Application ID** – The ID assigned to your application.
 - **Subtitle** – The subtitle that you chose to assign to your application.
 - **Service access** – The service access role that your application is using.
 - **Title** – The title that you gave to your application.
 - **Application status** – The status of your application.

To update a setting, select **Edit**.

AWS CLI

To get Amazon Q application properties

```
aws qbusiness get-application \  
--application-id application-id
```

Listing applications

To list Amazon Q Business applications, you can use the console or the [ListApplications](#) API operation.

The following tabs provide a procedure for the console and code examples for the AWS CLI.

Console

To list your Amazon Q applications

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, all your configured applications are listed.

AWS CLI

To list Amazon Q applications

```
aws qbusiness list-applications \  
--max-results max-results-to-return
```

Updating an application

To update an Amazon Q Business application, you can use the console or the [UpdateApplication](#) API operation.

Note

You can't update the retriever you've chosen or change users and groups added to the application when you update it. If you need to update your retriever, create a new application.

If you're integrating your Amazon Q application with IAM Identity Center (IDC) as an [AWS-managed](#) application using and you want to update users and groups, you can do so from the [application summary](#) page.

The following tabs provide a procedure for the console and code examples for the AWS CLI.

Console

To update an Amazon Q application

Option 1

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, select the name of your application from the list of applications.
3. In **Applications**, choose **Actions**.
4. Choose **Edit**.

On the **Update application** page, edit your application settings.

Option 2

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, select the name of your application from the list of applications.
3. On your application page, select **Edit** from the page header, or select **Edit** from **Application settings**.
4. Choose **Edit**.

On the **Update application** page, edit your application settings.

AWS CLI

To update an Amazon Q application

```
aws qbusiness update-application \  
--application-id application-id \  
--display-name application-name \  
--role-arn roleArn \  
--description application-description \  
--attachments-configuration attachmentsControlMode=ENABLED
```

Creating and selecting a retriever for an Amazon Q Business application

After creating your Amazon Q Business application, you create and select the retriever that will power your generative AI web experience. A retriever pulls data from an index in real time during a conversation. Amazon Q provides retrievers for Amazon Kendra indexes and also for a native index. You can choose between selecting an Amazon Q retriever or using an already configured Amazon Kendra index as a retriever.

To select a retriever, you use the AWS Management Console or the [CreateRetriever](#) API operation.

If you use the console and choose to use a Amazon Q retriever, Amazon Q creates an index for you as part of the application configuration process. For easy tracking, you can tag both the retriever and index. If you use the API to create a Amazon Q retriever, you must also use the [CreateIndex](#) API operation to create an Amazon Q index.

⚠ Important

You can't change the retriever for your application after your application has been created. To change your retriever, you must create a new application.

ℹ Note

The data sources available to connect to your application change depending on your retriever choice.

For instructions on how to select a retriever, choose a topic based on your retriever preference for Amazon Q.

Topics

- [Creating an Amazon Q Business retriever](#)
- [Selecting an Amazon Kendra retriever to an Amazon Q Business application](#)

Creating an Amazon Q Business retriever

To select a Amazon Q Business retriever, you can use either the AWS Management Console, or the [CreateIndex](#) and [CreateRetriever](#) API operations.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console**To create an Amazon Q retriever**

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Then, for **Select retriever**, choose **Use native retriever** – Build an Amazon Q Business retriever for your Amazon Q Business application. This option creates an Amazon Q Business index that can connect to the Amazon Q Business supported data sources that you choose.

⚠ Important

The native retriever includes a default capacity of 10k documents and 0.5 queries per second (QPS).

ℹ Note

Available data sources when you select this option include all [Amazon Q supported data connectors](#) and direct document upload.

4. For **Index provisioning** – Choose the **Number of units** that you need. Amazon Q charges you based on the document capacity that you choose. You can choose up to 50 units. Each unit is 20,000 documents or 200 MB, whichever comes first.
5. For **Tags** – Choose whether you want to add **Index tags**.
6. To create your retriever, choose **Create**.

AWS CLI**To create an Amazon Q index**

```
aws qbusiness create-index \  
--application-id application-id \  
--display-name display-name \  
--description index-description \  
--capacity-configuration units =<index-capacity-units>
```

To create an Amazon Q retriever

```
aws qbusiness create-retriever \  
--application-id application-id \  
--display-name display-name \  
--type NATIVE_INDEX \  
--role-arn roleArn \  
--configuration nativeIndexConfiguration="{indexId=<created-index-id>}" \  
--tags tags
```


Managing Amazon Q Business retrievers

To manage Amazon Q Business retrievers, you can take the following actions:

Actions

- [Deleting an Amazon Q Business retriever](#)
- [Getting properties of an Amazon Q Business retriever](#)
- [Listing Amazon Q Business retrievers](#)
- [Updating Amazon Q Business retrievers](#)

Deleting an Amazon Q Business retriever

To delete a Amazon Q Business retriever and its associated index, you can use the console or the [DeleteRetriever](#) API operation.

If you use the `DeleteIndex` API operation, deleting a retriever also deletes the Amazon Q index that's attached to it. You can't selectively choose to delete an index attached to a retriever.

If you're using the console, the only way to delete your Amazon Q native retriever and the index associated with it, is to delete your Amazon Q application.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To delete an Amazon Q retriever

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, choose **Actions**.
3. Choose **Delete**.
4. In the dialog box that opens, type **Delete** to confirm deletion, and then choose **Delete**.

You are returned to the service console while your application is deleted. When the deletion process is complete, the console displays a message confirming successful deletion.

AWS CLI

To delete an Amazon Q retriever

```
aws qbusiness delete-retriever \  
--application-id application-id \  
--retriever-id retriever-id
```

Getting properties of an Amazon Q Business retriever

To get the properties of an Amazon Q Business retriever and index, you can use the console or the [GetRetriever](#) API operation.

Note

If you use the console, you can't edit or update retriever or index settings.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To get properties of an Amazon Q retriever

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, select the name of your application from the list of applications.
3. For **Retriever settings**, the following settings are available:
 - **Retriever** – The type of retriever that you're using.
 - **Document count** – The number of documents that are attached to your index.
 - **Last modified time** – The time that your index was last modified.
 - **Index ID** – The ID of the index attached to your retriever.
 - **Storage used** – The amount of storage that your index is using.
 - **Index status** – The status of your index.

AWS CLI

To get properties of an Amazon Q retriever

```
aws qbusiness get-retriever \  
--application-id application-id \  
--retriever-id retriever-id
```

Listing Amazon Q Business retrievers

To list your native Amazon Q Business retrievers, you can use the console or the [ListRetrievers](#) API operation.

If you use the console, the list of Amazon Q retrievers and indices attached to them correspond to the list of applications that you have created.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To list your Amazon Q retrievers

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. For **Applications**, a list of all retrievers (with indices associated) that you have created is available.

AWS CLI

To list your Amazon Q retrievers

```
aws qbusiness list-retrievers \  
--application-id application-id \  
--max-results maximum-result-to-display
```

Updating Amazon Q Business retrievers

To update your Amazon Q Business retriever, you can use the [UpdateRetriever](#) API operation.

You can't update your retriever and its associated index by using the console.

The following tab provides code examples for the AWS CLI.

Console

This action is not supported on the console.

AWS CLI

To update your Amazon Q retriever

```
aws qbusiness update-retriever \  
--application-id application-id \  
--retriever-id retriever-id \  
--display-name display-name \  
--role-arn roleArn \  
--configuration kendraIndexConfiguration="{indexId=<kendra-index-id>}"
```

Selecting an Amazon Kendra retriever to an Amazon Q Business application

To select an existing Amazon Kendra retriever to your Amazon Q Business application, you can use the AWS Management Console or the [CreateRetriever](#) API operation.

If you use the API, you select and connect your Amazon Kendra retriever when you use the `CreateRetriever` API operation.

If you use the console, selecting and connecting an Amazon Kendra retriever is a two-step process. This topic provides instructions for the first step: Selecting an Amazon Kendra retriever. For instructions for the second step, see [Connecting an Amazon Kendra retriever to an Amazon Q Business application](#).

Note

If you use an Amazon Kendra retriever, data in your Amazon Kendra will be connected to your Amazon Q application. If you choose this option, you can't use Amazon Q data connectors or direct document upload for your application.

For more information about Amazon Kendra, see the following topics in the Amazon Kendra User Guide and API Reference:

- [What is Amazon Kendra?](#)
- [Creating a data source connector](#)
- [Amazon Kendra API Reference](#)

The following tabs provide a procedure for the AWS Management Console and code samples for the AWS CLI.

Console

To create an Amazon Kendra retriever

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. The, in **Select retriever** choose **Use existing retriever** – Choose an Amazon Kendra index you have previously created as a retriever. All data sources synced to your Amazon Kendra index will be connected to your Amazon Q Business application.
4. In **Tags** – Choose whether you want to add **Retriever tags**.
5. To connect your application to your data sources, choose **Next**.

AWS CLI

To create an Amazon Kendra retriever

```
aws qbusiness create-retriever \  
--display-name display-name \  

```

```
--type KENDRA_INDEX \  
--role-arn roleArn \  
--configuration kendraIndexConfiguration="{indexId=<kendra-index-id>
```

Managing Amazon Kendra retrievers

To manage Amazon Kendra retrievers, you can take the following actions:

Actions

- [Deleting an Amazon Kendra retrievers](#)
- [Getting properties of an Amazon Kendra retriever](#)
- [Listing Amazon Kendra retrievers](#)
- [Updating an Amazon Kendra retriever](#)

Deleting an Amazon Kendra retrievers

To delete an Amazon Kendra retriever, you can use the console or the [DeleteRetriever](#) API operation.

If you use the console, the only way to delete your Amazon Kendra retriever from your Amazon Q Business application is to delete your Amazon Q Business application.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To delete an Amazon Kendra retriever

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, choose **Actions**.
3. Choose **Delete**.
4. In the dialog box that opens, type **Delete** to confirm deletion, and then choose **Delete**.

You are returned to the service console while your application is deleted. When the deletion process is complete, the console displays a message confirming successful deletion.

AWS CLI

To delete an Amazon Kendra retriever

```
aws qbusiness delete-retriever \  
--application-id application-id \  
--retriever-id retriever-id
```

Getting properties of an Amazon Kendra retriever

To get the properties of an Amazon Kendra retriever, you can use the console or the [GetRetriever](#) API operation.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To get the properties of an Amazon Kendra retriever

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, select the name of your application from the list of applications.
3. For **Retriever settings**, the following settings are available:
 - **Retriever** – The type of retriever that you're using.
 - **Document count** – The number of documents that are attached to your index.
 - **Last modified time** – The time that your index was last modified.
 - **Index ID** – The ID of the index attached to your retriever.
 - **Storage used** – The amount of storage that your index is using.
 - **Index status** – The status of your index.

Note

You can't edit or update retriever or index settings.

AWS CLI

To get properties of an Amazon Kendra retriever

```
aws qbusiness get-retriever \  
--application-id application-id \  
--retriever-id retriever-id
```

Listing Amazon Kendra retrievers

To list Amazon Kendra retrievers, you can use the console or the [ListRetrievers](#) API operation.

If you use the console, the list of native retrievers and indices attached to them correspond to the list of applications that you have created.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To list Amazon Kendra retrievers

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. For **Applications**, a list of all retrievers (with indices associated) that you have created is available.

AWS CLI

To list Amazon Kendra retrievers

```
aws qbusiness list-retrievers \  
--application-id application-id \  
--max-results maximum-result-to-display
```

Updating an Amazon Kendra retriever

To update your Amazon Kendra retriever, you can use the [UpdateRetriever](#) API operation.

You can't update your Amazon Kendra retriever using the console.

The following tab provides code examples for the AWS CLI.

Console

This action is not supported on the console.

AWS CLI

To update an Amazon Kendra retriever

```
aws qbusiness update-retriever \  
--application-id application-id \  
--retriever-id retriever-id \  
--display-name display-name \  
--role-arn roleArn \  
--configuration kendraIndexConfiguration="{indexId=<kendra-index-d>}"
```

Connecting data sources to an Amazon Q Business application

After you select a retriever for your Amazon Q Business application, you connect data sources to it. Available data sources vary based on your choice of the retriever.

If you use an Amazon Q retriever, you can choose from the following options:

- Connect to any Amazon Q supported data source connectors by using the [CreateDataSource](#) API operation.
- Upload documents directly by using the [BatchPutDocument](#) API operation.

If you use an existing Amazon Kendra retriever, only data sources already connected to your Amazon Kendra index are available in your application.

To connect data sources, choose a topic based on your data source preference for your Amazon Q application.

Topics

- [Upload documents](#)

- [Amazon Kendra retriever](#)
- [Amazon Q data source connectors](#)

Upload documents

To upload documents directly to an Amazon Q Business application, you can use the AWS Management Console or the [BatchPutDocument](#) API operation.

If you use an Amazon Kendra index to retrieve your documents, you can't directly upload documents.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To upload documents

Note

This procedure is available if you chose the **Use native retriever** option to configure your application.

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, for **Upload documents**, select one of the following methods to add your files:
 - Drag and drop the document files that you want to upload.
 - Add your documents to the application, and then select **Choose files**.
5. After choosing your files, choose **Upload**.

You are returned to the Amazon Q console while your documents are uploaded. The console displays a confirmation message when your documents are successfully uploaded.

Note

Files can only be uploaded after the Amazon Q retriever and index creation process has completed.

AWS CLI

To upload documents directly

```
aws qbusiness batch-put-document \  
--application-id application-id \  
--index-id index-id \  
--documents documents-to-add \  
--data-source-sync-id data-source-sync-id \  
--role-arn roleArn
```

Delete uploaded documents

To delete documents that have been directly uploaded to an application, you can use the console or the [BatchDeleteDocument](#) API operation. You can delete specific documents or all documents.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To delete specific directly uploaded documents

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, select the name of the application that your uploaded files belong to.
3. From your applications page, from **Data sources**, choose **Uploaded files**.
4. In **Uploaded files**, choose **Document name**, and then select the documents that you want to delete.
5. Choose **Delete files**.

You are returned to the service console while your application is deleted. When the deletion process is complete, the console displays a message confirming successful deletion.

To delete all directly uploaded documents

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, select the name of the application that your uploaded files belong to.
3. From your applications page, from **Data sources**, select **Uploaded files**.
4. Select **Actions**, and then choose **Delete**.
5. When the deletion process is complete, the console displays a message confirming successful file deletion.

AWS CLI

To delete documents

```
aws qbusiness batch-delete-document \  
--application-id application-id \  
--index-id index-id \  
--documents documents-to-delete \  
--data-source-sync-id data-source-sync-id
```

Connecting an Amazon Kendra retriever to an Amazon Q Business application

To use an Amazon Kendra index as a retriever for Amazon Q Business, you must have already configured an Amazon Kendra index and connected it with data. For more information, see [What is Amazon Kendra?](#) and [Are you a first-time Amazon Kendra user?](#) in the Amazon Kendra Developer Guide.

To add an existing Amazon Kendra retriever to your Amazon Q application, you can use the AWS Management Console or the [CreateRetriever](#) API operation. If you use the console, selecting and connecting an Amazon Kendra retriever is a two-step process. The first step is when you [select](#)

[an Amazon Kendra retriever](#). In this topic, you perform the second step—connecting an Amazon Kendra retriever.

If you use the API, you create your web experience after connecting your Amazon Kendra retriever using the [CreateWebExperience](#) API operation. If you use the console, connecting your Amazon Kendra retriever also automatically creates your Amazon Q web experience. At the end of the retriever connection process, your Amazon Kendra powered Amazon Q web experience is ready to be previewed, enhanced, and deployed.

Note

If you select an Amazon Kendra retriever, data in your Amazon Kendra is connected to your Amazon Q application.

Console

To connect an Amazon Kendra retriever

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Kendra retriever](#).
4. Then, in **Content sources**, for **Amazon Kendra indexes** – Choose the Amazon Kendra index that you want to use for your Amazon Q Business application. Then, enter the following information:
 - **Service access** – Provide the IAM access role to connect Amazon Kendra to Amazon Q Business. Use an existing role, or create a new one.
 - **Service role name** – Provide a name for your IAM access role. Or, choose to use the auto-generated role that's provided.
5. To connect your Amazon Kendra indexes to the application, choose **Create application**.

You are returned to the Amazon Q console while your web application is created.

AWS CLI

To create and connect an Amazon Kendra retriever

```
aws qbusiness create-retriever \  
--application-id application-id \  
--display-name display-name \  
--type KENDRA_INDEX \  
--role-arn roleArn \  
--configuration kendraIndexConfiguration="{indexId=<kendra-index-id>}"
```

Note

For information on managing your Amazon Kendra retriever, see [Managing Amazon Kendra retrievers](#).

Amazon Q Business data sources

To connect a data source to your Amazon Q Business application, you can use the AWS Management Console or the [CreateDataSource](#) API operation.

By using the `CreateDataSource` API operation, you can configure tags, sync run schedules, and configure Amazon VPC settings. Then, you can use the `configuration` parameter to provide all other configuration information specific to your data source connector.

If you use the console, creating the data source and configuring it are a single step. After your data source is successfully configured and added, Amazon Q automatically creates a Amazon Q web experience for you.

If you use the API, you use the [CreateWebExperience](#) API operation after connecting your data sources to create your web experience.

Note

This procedure is available if you chose the [Use native retriever](#) option to configure your application.

Console

To connect a data source to an Amazon Q application

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 50 data sources.

5. For information on configuring your chosen data source, see [Supported connectors](#) to find configuration information specific to your data source.
6. To connect your configured data source to your application, choose **Add data sources**.

At the end of this step, your Amazon Q web experience is ready to be previewed, enhanced, and deployed.

AWS CLI

To connect a data source

```
aws qbusiness create-data-source \
--application-id application-id \
--index-id index-id \
--configuration data-source-configuration-details \
--display-name display-name \
--role-arn roleArn \
--description description \
--document-enrichment-configuration document-enrichment-configuration \
--sync-schedule sync-schedule-information \
--tags tags \
--vpc-configuration vpc-configuration
```

Managing Amazon Q Business data sources

To manage data source connectors, you can perform the following actions:

Actions

- [Deleting an Amazon Q Business data source connector](#)

- [Getting properties of an Amazon Q Business data source connector](#)
- [Listing Amazon Q Business data source connectors](#)
- [Updating Amazon Q Business data source connectors](#)
- [Starting data source connector sync jobs](#)
- [Stopping data source connector sync jobs](#)
- [Listing data source connector sync jobs](#)

Deleting an Amazon Q Business data source connector

To delete an Amazon Q Business data source connector, you can use the console or the [DeleteDataSource](#) API operation .

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To delete an Amazon Q data source connector

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, select the application that you want to delete data sources from.
3. On the application page, from **Data sources**, select the data source that you want to delete.
4. From **Actions**, choose **Delete**.
5. In the dialog box that opens, type **Delete** to confirm deletion, and then choose **Delete**.

You are returned to the service console while your data source connector is deleted. When the deletion process is complete, the console displays a message confirming successful deletion.

AWS CLI

To delete an Amazon Q data source connector

```
aws qbusiness delete-data-source \  
--application-id application-id \  
--index-id index-id \  

```



```
--data-source-id data-source-id
```

Getting properties of an Amazon Q Business data source connector

To get the properties of an Amazon Q Business data source connector, you can use the [GetDataSource](#) API operation.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To get properties of an Amazon Q data source connector

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, select the application you want that contains your data sources.
3. On the application page, from **Data sources**, select the data source that you want to view details for.
4. Under **Data source details**, the following details are available:
 - **Name** – The name of your data source.
 - **Status** – The status of your data source.
 - **Last sync status** – The status of your last sync.
 - **Description** – The description that you gave to your data source.
 - **Type** – The type of data source that you're using.
 - **Last sync time** – The time that your data source was last synced.
 - **Data source ID** – The ID of your data source.
 - **IAM role ARN** – The Amazon Resource Name (ARN) of the IAM role that's associated with your data source.
 - **Current sync state** – The current sync state of your data source.

To get Amazon Q data source connector settings

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.

2. In **Applications**, select the application you want that contains your data sources.
3. On the application page, from **Data sources**, select the data source that you want to view details for.
4. For **Data source details**, choose **Settings**.
5. For **Settings**, the following settings are available:
 - **IAM role** – The ARN of the IAM that's associated with your data source.
 - **Sync scope** – The configuration details for your data source.
 - **Sync mode** – The sync type that you chose for your data source.
 - **Sync schedule** – The sync schedule that you chose for your data source.
 - **Field mappings** – The data source document fields that you chose to map to Amazon Q index fields.

AWS CLI

To get Amazon Q data source connector properties

```
aws qbusiness get-data-source \  
--application-id application-id \  
--index-id index-id \  
--data-source-id data-source-id
```

Listing Amazon Q Business data source connectors

To list Amazon Q Business data source connectors, you can use the console or the [ListDataSources](#) API operation.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To list Amazon Q data source connectors

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, select the application you want that contains your data sources.

3. On the application page, under **Data sources**, a list of data sources connected to your application is displayed.

AWS CLI

To list Amazon Q data source connectors

```
aws qbusiness list-data-sources \  
--application-id application-id \  
--index-id index-id \  
--max-results maximum-number-of-results-to-return
```

Updating Amazon Q Business data source connectors

To update your Amazon Q Business data source connectors, you can use the console or the [UpdateDataSource](#) API operation.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To update a Amazon Q data source connector

Option 1

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, select the application you want to delete data sources from.
3. On the application page, from **Data sources**, select the data source that you want to edit.
4. From **Actions**, choose **Edit**.

You are redirected to your data source configuration page to edit your existing settings.

Option 2

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.

2. In **Applications**, select the application you want to delete data sources from.
3. On the application page, from **Data sources**, select the data source that you want to edit.
4. On the data source page, from **Actions**, choose **Edit**.

You are redirected to your data source configuration page to edit your existing settings.

CLI

To update your Amazon Q connector

```
aws qbusiness update-data-source \  
--application-id application-id \  
--data-source-id data-source-id \  
--index-id index-id \  
--configuration data-source-configuration-details \  
--description description \  
--display-name display-name \  
--document-enrichment-configuration document-enrichment-configuration \  
--role-arn roleArn \  
--sync-schedule sync-schedule-information \  
--vpc-configuration vpc-configuration
```

Starting data source connector sync jobs

To start Amazon Q Business data source connector sync jobs, you can use the console or the [StartDataSourceSyncJobs](#) API operation.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To start your Amazon Q data source connector sync jobs

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/aamazonq/>.
2. In **Applications**, select the application you want to sync data sources in.
3. On the application page, from **Data sources**, select the data source that you want to sync.

4. Choose **Sync now**.

The console displays a message confirming that your sync job has started successfully.

Note

You can also view your sync job report in the Amazon CloudWatch console.

AWS CLI

To start your Amazon Q data source connector sync jobs

```
aws qbusiness start-data-source-sync-job \  
--application-id application-id \  
--index-id index-id \  
--data-source-id data-source-id
```

Stopping data source connector sync jobs

To stop your Amazon Q Business connector sync jobs, you can use the console or the [StopDataSourceSyncJobs](#) API operation.

Note

You can only stop a sync job already in progress.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To stop your Amazon Q data source connector sync jobs

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, select the application you want to sync data sources in.

3. On the application page, from **Data sources**, select the data source that you want to stop the sync for.
4. Choose **Stop sync**.
5. In the dialog box that opens, type **Stop** to confirm your action and then select **Stop sync**.

The console displays a message confirming that your data source sync job is being stopped.

AWS CLI

To stop your Amazon Q data source connector sync jobs

```
aws qbusiness stop-data-source-sync-job \  
--application-id application-id \  
--data-source-id data-source-id \  
--index-id index-id
```

Listing data source connector sync jobs

To list Amazon Q Business data source connector sync jobs that are in progress, you can use the console or the [ListDataSourceSyncJobs](#) API operation.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To list your Amazon Q data source connector sync jobs

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, select the application you want that contains your data sources.
3. On the application page, from **Data sources**, select the data source that you want to view details for.
4. Under **Data source details**, choose the **Sync run history** tab.

You will see a list of ongoing, completed, and failed sync jobs for your data sources.

CLI

To list your Amazon Q data source connector sync jobs

```
aws qbusiness list-data-source-sync-job \  
--application-id application-id \  
--data-source-id data-source-id \  
--index-id index-id \  
--max-results max-results-to-return
```

Adding users and groups to an Amazon Q application

For your end users to log in and chat, Amazon Q requires that you add at least one user to an IAM Identity Center integrated application. This section outlines the steps for this process. This integration is required so that only authorized end users from within your organization have access to your content.

To create your Amazon Q web experience, you can use either the AWS Management Console or the Amazon Q API. If you choose the API, use the [CreateWebExperience](#) API operation to create and deploy your web experience.

If you use the console to create your Amazon Q application, a web experience is created automatically.

On successful completion, Amazon Q Business returns a web experience URL that you can share with the end users you added to your application.

Note

During Preview, an Amazon Q application supports only 50 end users. If you need more capacity, contact [Support](#).

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To add users and groups to your Amazon Q application

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. (Optional) Complete the steps for [connecting data sources](#).
5. To add users, in **Add groups and users**, choose **Users**, and then choose **Add users**. Then, depending on how you're integrating Amazon Q with IAM Identity Center, do the following:
 - a. If you're using a pre-configured IAM Identity Center instance with users and groups already added, Amazon Q detects the users you have configured in IAM Identity Center. You can choose to assign users from your IAM Identity Center directory.
 - i. In this case, in the **Assign users and groups** dialog box that opens, type and select the name of the users you want to assign.
 - ii. Then, select the user you want to add and then select **Assign**.
 - b. If you've created a minimally-configured IAM Identity Center instance from within the Amazon Q console for your Amazon Q application, you can add users to your IAM Identity Center instance and then assign new users to your Amazon Q application in this step.
 - i. In this case, in the **Add or assign users and groups** dialog box that opens, select **Add new users** and select **Next**.
 - ii. In **Add new users**, select **Add new users**.
 - iii. In **Add new users**, enter the following information:
 - **Username** – A username is required for an user to sign into the AWS access portal. You can't change the username later. Maximum length 128 characters. Can only contain alphanumeric characters or any of the following: +=,.,@-_
 - **First name** – First name of user.
 - **Last name** – Last name of user.
 - **Email address** – Email address of user.
 - **Confirm email address** – Enter email address again to confirm it.
 - **Display name** – The display name assigned to your user.
 - iv. Then, select **Next**. The **Add new users** dialog box displays the user who you've just added.

v. **Select Add.**

The new user you add will receive a notification to the email you entered for them asking them to accept your invitation to IAM Identity Center. You will also have to go to the IAM Identity Center console and send them an email verification request. Your user will have to complete the email verification before they can successfully log in to the web experience URL for your Amazon Q application. For more information on enabling users in IAM Identity Center, see [Add users](#) in the IAM Identity Center User Guide.

vi. From the **Add or assign users and groups** dialog box that opens, choose **Assign existing users and groups** and enter and select the name of the user you added in the last step. Then, choose **Assign**.

When the assignment process successfully completes, the Amazon Q console will display the name, username, and user ID of the user assigned to your Amazon Q application.

6. To add groups, in **Add groups and users**, choose **Groups**, and then choose **Add groups**. Then, depending on how you're integrating Amazon Q with IAM Identity Center, do the following:

- a. If you're using a pre-configured IAM Identity Center instance with users and groups already added, Amazon Q detects the groups you have configured in IAM Identity Center. You can choose to assign groups from your IAM Identity Center directory.
 - i. In this case, in the **Assign users and groups** dialog box that opens, type and select the name of the groups you want to assign.
 - ii. Then, select the group you want to add and select **Assign**.

 **Note**

As an IAM Identity Center user, you also have the option to add and adding a new user to your app from within the Amazon Q console. See the following step for details on how to do so. However, you can't create groups from within the Amazon Q console.

- b. If you've created a minimally-configured IAM Identity Center instance from within the Amazon Q console for your Amazon Q application, you can't add groups to your IAM Identity Center instance from the Amazon Q console.
 - i. In this case, in the **Add or assign users and groups** dialog box that opens, select **Manage users and groups from IAM Identity Center**.
 - ii. From the left navigation menu, choose **Groups**.
 - iii. From **Groups**, choose **Create group**. Then, enter the following information:
 - **Group name** – A name for your group.
 - **Description** – An optional group description.
 - iv. In **Add users to groups**, add users you want to the group.

 **Note**

The users must already exist in your IAM Identity Center instance. If they don't, you can select **Add new user** and follow the steps to add a new user to the group you're creating.

- v. Then, select **Create group**. The console will display a group summary page.
7. In **Web experience service access**, enter the following information:
 - For **Choose a method to authorize Amazon Q** – A service access role assumed by end users when they sign in to your web experience that grants them permission to start and manage conversations Amazon Q. You can choose to use an existing role or create a new role.
 - **Service role name** – A name for the service role you created for easy identification on the console.
 8. Select **Save**.
 9. Select **Create application**.

AWS CLI

To add users to an application

```
aws sso-admin create-application-assignment \
```

```
--application-arn idc-app-arn \  
--principal-id idc-user-ID \  
--principal-type USER
```

To add groups to an application

```
aws sso-admin create-application-assignment \  
--application-arn idc-app-arn \  
--principal-id idc-group-ID \  
--principal-type GROUP
```

Configuring Amazon Q Business data source connectors

A *data source connector* is a mechanism for integrating and synchronizing data from multiple repositories into one container index. Amazon Q Business offers multiple data source connectors that can connect to your data sources and help you create your generative AI solution with minimal configuration.

To configure and connect a data source to your Amazon Q application, use the [CreateDataSource](#) API operation. Specify your connector configuration details using the `configuration` parameter of the `CreateDataSource` operation. If you use the AWS Management Console instead of the API, you create, configure, and connect your data source as part of the application creation process.

This section contains an overview of data source connector features, recommended best practices for configuration, and configuration information specific to your data source connector.

Topics

- [Data source connector concepts](#)
- [Best practices for data source connector configuration in Amazon Q Business](#)
- [Supported connectors](#)
- [Understanding Amazon Q Business User Store](#)
- [Using Amazon VPC with Amazon Q Business connectors](#)
- [Troubleshooting data source connectors](#)

Data source connector concepts

This topic outlines specific concepts and features of Amazon Q Business data source connectors. These concepts are key to understanding how to configure your connector setup. These terms recur on the AWS Management Console, AWS Command Line Interface (AWS CLI), and the Amazon Q API.

Topics

- [Source and endpoint metadata](#)
- [Authorization](#)
- [Authentication](#)

- [Virtual private cloud](#)
- [Web proxy](#)
- [IAM role](#)
- [Identity crawler](#)
- [Sync scope](#)
- [Sync mode](#)
- [Sync run schedule](#)
- [Field mappings](#)

Source and endpoint metadata

You enter your data source configuration information in the **Source** section on the console. If you use the API, you specify this information using the `configuration` parameter of the `CreateDataSource` operation. Connection configuration information varies depending on the data source. To make sure your connector configures correctly, check the following details:

- You're following [connector configuration best practices](#).
- You've completed the prerequisites for data source configuration. Prerequisites information specific to your data source connector is on each connector's specific page.

Authorization

Amazon Q Business connectors crawl access control list (ACL) information that's attached to a document along with the document itself. Specifically, the crawled ACL information is the local users and local groups that have access to the document.

Then, Amazon Q uses the ACL information it crawls, together with the [Identity crawler](#) feature, to generate chat responses for your end user. The responses are based on the documents that the user has access to. For more information about this process, see [Understanding User Store](#).

During connector configuration, you can choose either to activate or deactivate the **Authorization** feature. If you activate this feature, Amazon Q will only generate content for an end user from documents they have access to. If this feature is deactivated, the connector won't crawl any ACL information and all documents will be considered public to all application users.

If you choose to crawl an ACL, the connector updates any changes in ACLs each time that your data source content is crawled. To capture ACL changes to make sure that the right end users have access to the right content, re-sync your data source regularly.

Authentication

To authenticate Amazon Q Business to access your data source, you provide your data source access credentials to Amazon Q using an AWS Secrets Manager secret. If you use the console, you can choose to create a new secret or use an existing one. If you use the API, you must provide the Amazon Resource Name (ARN) of an existing Secrets Manager secret when you use the `CreateDataSource` operation.

Note

You should regularly refresh or rotate your credentials and secret details. Provide only the necessary access level for your own security. Don't re-use credentials and secrets across data sources.

For on-premises or server data source connectors, Amazon Q checks if the endpoint information included in Secrets Manager is the same endpoint information specified in your data source configuration details. This helps protect against the [confused deputy problem](#), which is a security issue. The problem occurs when a user doesn't have permissions to perform an action. But, by using Amazon Q as a proxy, the user can access the configured secret and perform the action.

If you change your endpoint information later, you must create a new secret to sync this information.

Note

If you change your authentication type and credentials, you must update your IAM role to access the correct Secrets Manager secret ID.

Virtual private cloud

Amazon Q Business can connect to Amazon Virtual Private Cloud to index content stored in data sources or databases running in your private cloud. If your data source or database isn't running on

Amazon VPC, you can connect your data source or database to Amazon VPC using a virtual private network (VPN).

You can use Amazon VPC with either the console or the Amazon Q API. If you're using the API, you specify the `vpcConfiguration` when you use the `CreateDataSource` API operation.

If you're using Amazon VPC with Amazon Q, you need the following information:

- The identifier of the subnet that contains the data source.
- The identifier of the security groups that grant access to the host.
- An IAM role with access to Amazon VPC and permissions to create and delete an elastic network interface in your subnets is also required.

You can find the subnet and security group IDs in the Amazon VPC console. For more information, see [What is Amazon VPC?](#) in the *Amazon VPC User Guide*.

For more information about using Amazon VPC with Amazon Q, see [Using Amazon VPC with connectors](#).

Web proxy

For all supported data sources, you can use a web proxy to connect to your data source instance. You must provide the host name and port number. For example, `a.example.com` is the hostname of `https://a.example.com/page1.html`, and the port is 443, which is the standard port for HTTPS.

Important

For security reasons, Amazon Q Business only supports web proxy using HTTPS protocol.

IAM role

To create your data source connector, Amazon Q Business requires permissions to interact with other services.

If you're using the console, you can choose an existing IAM role or let Amazon Q create a role for you. If you're unsure if an existing role is used for an application, choose **Create a new role** to avoid an error.

Note

To **Create a new role** during connector configuration on the console, you must have permissions to create an IAM role.

If you're using the API, you must provide the ARN of an existing IAM role when you use the `CreateDataSource` operation.

IAM roles used for applications can't be used for data sources.

Note

Make sure your IAM role includes the permissions to support your Amazon Q connector configurations.

Identity crawler

Amazon Q Business crawls ACL information at the document level from supported data sources (users and group access to documents). In addition, Amazon Q crawls the local user and group identity configurations within each data source. This approach is useful where your application is connected to multiple data sources with different authorization and authentication systems, but you want to create a unified, access-controlled chat experience for end users.

Amazon Q stores the principal information in the Amazon Q user store. Then, Amazon Q internally maps the local user and group IDs attached to the document to the federated identities of users and groups. Mapping identities streamlines user management and speeds up chat responses by reducing ACL information retrieval time during chat requests.

Identity crawling, along with the [Authorization](#) feature, helps to filter and generate web experience content restricted by end user context. For more information about this process, see [Understanding User Store](#).

Note

To activate the **Identity crawler** feature on the console, you must also activate ACL crawling as part of the **Authorization** feature. If you use the API, you choose to activate

identity crawling using the configuration object when you use the `CreateDataSource` operation.

Sync scope

You can choose to customize the content crawled and indexed by your data source connector. The sync scope options available vary based on the data source connector.

Sync mode

With sync mode, you can customize what content gets synced with your index when your data source content changes. Choose from the following options:

Console

- **Full sync** – Sync all content regardless of the previous sync status.
- **New or modified content sync** – Sync only new or modified documents.
- **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.
- **Change log** – Crawl and sync only new, modified, and deleted content.

API

Specify the sync mode using the `configuration` parameter of the [CreateDataSource](#) operation. Choose from the following options:

- **Forced full crawl** – Crawl and sync all content to your index.
- **Full crawl** – Crawl all content and sync only new, modified, or deleted content.
- **Change log** – Crawl and sync only new, modified, and deleted content.

Note

Available sync mode features vary across data source connectors.

Important

If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it's because the CloudWatch logs aren't available yet. Wait for some time and check again.

Sync run schedule

When you use the console or the [CreateDataSource](#) API operation, you can choose to periodically sync your data source with your retriever on a custom schedule. You can choose from the following frequency options:

- **Run on demand** – Sync a data source with your index only when you choose to.
- **Hourly** – Sync your data source with your index every hour. You can choose which minute the sync begins.
- **Daily** – Sync your data source with your index daily. You can choose the sync start time in UTC format in hours and minutes.
- **Weekly** – Sync your data source with your index weekly. You can choose the days to sync and the sync start time in hours and minutes (UTC format).
- **Monthly** – Sync your data source monthly with your index. You can choose the day of the month to start the sync and the sync start time in hours and minutes (UTC format).
- **Custom** – Sync your data source to your index using a cron expression. A cron expression is a string comprising five or six required fields, separated by white space. Cron expressions represent a set of times programmed to schedule events. For example, an expression to activate a rule every day at 12:00pm UTC can look like: (0 12 * * ? *). Similarly, an expression to activate a rule every day at 10:15am UTC on the last Friday of each month during the years 2023 to 2025 can look like: (15 10 ? * 6L 2023-2025).

Note

Amazon Q will not sync the data source (even for the first time) until you select **Sync now** after you successfully add the data source.

Field mappings

When you connect Amazon Q Business to your data, your data source connector crawls relevant metadata or attributes associated with a document. Examples of metadata include date of creation, document id, and document name. Then, Amazon Q maps the metadata to fields within your Amazon Q index.

You map data source document attributes to Amazon Q index fields using the **Field mappings** feature on the console, or the `configuration` parameter of the `CreateDataSource` API operation. If you use the console, you add field mappings after your data source is created.

All fields and attributes have a size limit of 2048 characters. Fields or attributes longer than this value are truncated before document ingestion.

For more information, see the following topics:

- [Document attributes and types](#)
- [Filtering using metadata](#)

Best practices for data source connector configuration in Amazon Q Business

The following list describes best practices for setting up and configuring your Amazon Q Business data source connector:

- Each document in an index must be unique. Check that there are no duplicate documents in a data source, or across any data sources, that you plan to connect to an Amazon Q retriever.
- If you change your authentication type and credentials, you must update your IAM role to access the correct AWS Secrets Manager secret ID.
- We recommend that you regularly refresh or rotate your credentials and secret. Provide only the necessary access level for your own security. We don't recommend the re-use of credentials and secrets across data sources.
- IAM roles used for retrievers can't be used for data sources. If you're unsure if an existing role is used for a retriever or data source, create a new IAM role to avoid errors.
- If you use AWS KMS keys for the application, ensure that the IAM for your application is given the permission to describe, encrypt, and decrypt data using this key.

- For on-premises or server data source connectors, Amazon Q checks if the endpoint information included in Secrets Manager is the same as the endpoint information specified in your data source configuration details. This helps protect against the [confused deputy problem](#), which is a security issue. The problem occurs when a user doesn't have permission to perform an action. But, by using Amazon Q as a proxy, the user can access the configured secret and perform the action.

If you change your endpoint information later, you must create a new secret to sync this information.

- Most data sources use regular expression patterns, which are inclusion or exclusion patterns referred to as *filters*.

If you specify an inclusion filter, only content that matches the inclusion filter is indexed. If you specify an inclusion and exclusion filter, documents that match the exclusion filter aren't indexed, even if they match the inclusion filter.

Supported connectors

Amazon Q Business supports the following connectors:

- [AEM \(Cloud\)](#)
- [AEM \(Server\)](#)
- [Alfresco \(Cloud\)](#)
- [Alfresco \(Server\)](#)
- [Aurora \(MySQL\)](#)
- [Aurora \(PostgreSQL\)](#)
- [Amazon FSx Windows](#)
- [Amazon RDS \(Microsoft SQL Server\)](#)
- [Amazon RDS \(MySQL\)](#)
- [Amazon RDS \(Oracle\)](#)
- [Amazon RDS \(PostgreSQL\)](#)
- [Amazon S3](#)
- [Amazon Q custom data source connector](#)

- [Amazon Q Web Crawler](#)
- [Amazon WorkDocs](#)
- [Box](#)
- [Confluence \(Cloud\)](#)
- [Confluence \(Server\)](#)
- [Dropbox](#)
- [Drupal](#)
- [GitHub \(Cloud\)](#)
- [GitHub \(Server\)](#)
- [Gmail](#)
- [Google Drive](#)
- [IBM DB2](#)
- [Jira](#)
- [Microsoft Exchange](#)
- [Microsoft OneDrive](#)
- [Microsoft SharePoint \(Cloud\)](#)
- [Microsoft SharePoint \(Server\)](#)
- [Microsoft SQL Server](#)
- [Microsoft Teams](#)
- [Microsoft Yammer](#)
- [MySQL](#)
- [Oracle Database](#)
- [PostgreSQL](#)
- [Quip](#)
- [Salesforce Online](#)
- [ServiceNow Online](#)
- [Slack](#)
- [Zendesk](#)

Connecting AEM (Cloud) to Amazon Q Business

Adobe Experience Manager (AEM) is a content management system (CMS) that's used for creating website or mobile app content. You can connect your AEM (Cloud) instance to Amazon Q Business—using either the AWS Management Console, CLI, or the [CreateDataSource](#) API—and create an Amazon Q web experience.

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).
- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [AEM \(Cloud\) connector overview](#)
- [Prerequisites for connecting Amazon Q Business to AEM \(Cloud\)](#)
- [Connecting Amazon Q Business to AEM \(Cloud\) using the console](#)
- [Connecting Amazon Q Business to AEM \(Cloud\) using APIs](#)
- [How Amazon Q Business connector crawls AEM \(Cloud\) ACLs](#)
- [Amazon Q Business AEM \(Cloud\) data source connector field mappings](#)
- [IAM role for Amazon Q AEM \(Cloud\) connector](#)
- [Known limitations for the Amazon Q Business AEM \(Cloud\) connector](#)
- [Troubleshooting your Amazon Q Business AEM \(Cloud\) connector](#)

AEM (Cloud) connector overview

The following table gives an overview of the Amazon Q Business AEM (Cloud) connector and its supported features.

Category	Feature	Support
Security	Authentication type	Basic, OAuth 2.0

Category	Feature	Support
	Authentication credentials	<p>Basic</p> <ul style="list-style-type: none"> • AEM (Cloud) host URL • Username of AEM user • Password of AEM user <p>OAuth 2.0</p> <ul style="list-style-type: none"> • AEM (Cloud) host URL • Client ID • Client secret • Private key • Organization ID • Technical Account ID • Adobe Identity Management System (IMS) host <div style="border: 1px solid #f08080; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p>⚠ Important Admin privileges required.</p> </div>
	<u>Access Control List (ACL) crawling</u>	Yes. For more information, see <u>ACL crawling</u> .
	<u>Identity crawling</u>	Yes
	<u>VPC</u>	Yes
Crawl features	Custom metadata	Yes

Category	Feature	Support
	Entities	Yes. The following entities are supported: <ul style="list-style-type: none"> • Pages • Assets
	Field mappings	Yes. Supports both default and custom field mappings. For more information, see Field mappings .
	Filters	Yes. The following filters are supported: <ul style="list-style-type: none"> • Include/exclude by asset name • Include/exclude by asset type • Include/exclude by asset path • Include/exclude by page name • Include/exclude by page path
	Sync mode	Supports full and incremental sync.
	File types	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to AEM (Cloud)

Before you begin, make sure that you have completed the following prerequisites.

In AEM, make sure you have:

- Access to an account with administrative permissions, or are an admin user.
- Copied your AEM (Cloud) host URL.
- Noted your basic authentication credentials of admin username and password.
- (Optional) Added the following OAuth scopes if you're using OAuth 2.0 authentication:
 - **Profile** – Needed to get user and groups related data, like email ID and username.
 - **Replicate** – Needed to get data and metadata from Assets and Pages (not including user data).
- **Optional:** Generated OAuth 2.0 credentials in AEM (Cloud) as an admin user. The credentials include client ID, client secret, private key, organization ID, technical account ID, and Adobe

Identity Management System (IMS) host. For more information about how to generate these credentials for AEM (Cloud), see [AEM \(Cloud\) documentation](#).

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your AEM (Cloud) authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Connecting Amazon Q Business to AEM (Cloud) using the console

The following procedure outlines how to connect Amazon Q Business to AEM (Cloud) using the AWS Management Console.

Connecting Amazon Q to AEM (Cloud)

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **AEM (Cloud)** page, enter the following information:
6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. **Source** – Choose **AEM as a Cloud Service**.

- **AEM host URL** – Enter your **AEM host URL**. If you use AEM as a Cloud Service, you can use the author URL. For example: `https://author-xxxxx-xxxxxx-adobeemcloud.com`.

8. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

Note

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

9. **Authentication** – Choose between **Basic authentication** and **OAuth 2.0 authentication** and then enter the following information for your **AWS Secrets Manager secret**.

- a. **Basic authentication** – Enter a name for the secret, your AEM site admin username, and admin password.
- b. **OAuth 2.0 authentication** – Enter a name for the secret, your client ID, client secret, private key, organization ID, technical account ID, and Adobe IMS host.

10. **Configure VPC and security group** – *optional* – Choose whether you want to use a VPC. If you do, enter the following information:

- a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
- b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

11. **Identity crawler** – Choose to activate Amazon Q identity crawler to sync identity information. For more information, see [Identity crawler](#).

12. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

13. In **Sync scope**, enter the following information:
 - a. **Sync content types** – Choose whether to crawl only **Pages** or **Assets**, or both.
 - b. **Additional configuration – optional** – Configure the following settings:
 - **Page components** – The specific names of page components. The Page Component is an extensible page component designed to work with the Adobe AEM template editor and allows page header and footer and structure components to be assembled with the template editor.
 - **Content fragment variations** – The specific names of content fragment variations. Content Fragments allow you to design, create, curate, and publish page-independent content in Adobe AEM. They allow you to prepare content ready for use in multiple locations and over multiple channels.
 - **Root paths** – The root paths to specific content.
 - **Regex patterns** – The regular expression patterns to include or exclude certain pages and assets.
14. In **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.
 - **Full sync** – Sync all content regardless of the previous sync status.
 - **New or modified content sync** – Sync only new and modified documents.
 - **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.

For more details, see [Sync mode](#).
15. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
16. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
17. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:
 - a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.

- b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

 **Note**

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

18. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

19. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

 **Note**

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to AEM (Cloud) using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the `configuration` parameter to provide a JSON schema with all other configuration information specific to your data source connector.

For an example of the API request, see [CreateDataSource](#) in the Amazon Q API Reference.

AEM JSON schema

The following is the AEM JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
      "type": "object",
      "properties": {
        "repositoryEndpointMetadata": {
          "type": "object",
          "properties": {
            "aemUrl": {
              "type": "string",
              "pattern": "https:.*"
            },
            "authType": {
              "type": "string",
              "enum": [
                "Basic",
                "OAuth2"
              ]
            },
            "deploymentType": {
              "type": "string",
              "enum": [
                "CLOUD",
                "ON_PREMISE"
              ]
            }
          }
        },
        "required": [
          "aemUrl",
          "authType",
          "deploymentType"
        ]
      }
    },
    "required": [
      "repositoryEndpointMetadata"
    ],
    "repositoryConfigurations": {
      "type": "object",
      "properties": {
```

```
"page": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": [
        {
          "type": "object",
          "properties": {
            "indexFieldName": {
              "type": "string"
            },
            "indexFieldType": {
              "type": "string",
              "enum": [
                "STRING",
                "STRING_LIST",
                "DATE",
                "LONG"
              ]
            },
            "dataSourceFieldName": {
              "type": "string"
            },
            "dateFieldFormat": {
              "type": "string",
              "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
            }
          },
          "required": [
            "indexFieldName",
            "indexFieldType",
            "dataSourceFieldName"
          ]
        }
      ]
    },
    "required": [
      "fieldMappings"
    ]
  },
  "asset": {
    "type": "object",
```

```

    "properties": {
      "fieldMappings": {
        "type": "array",
        "items": [
          {
            "type": "object",
            "properties": {
              "indexFieldName": {
                "type": "string"
              },
              "indexFieldType": {
                "type": "string",
                "enum": [
                  "STRING",
                  "STRING_LIST",
                  "DATE",
                  "LONG"
                ]
              },
              "dataSourceFieldName": {
                "type": "string"
              },
              "dateFieldFormat": {
                "type": "string",
                "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
              }
            },
            "required": [
              "indexFieldName",
              "indexFieldType",
              "dataSourceFieldName"
            ]
          }
        ]
      },
      "required": [
        "fieldMappings"
      ]
    },
    "additionalProperties": {
      "type": "object",

```

```
"properties": {
  "isCrawlAcl": {
    "type": "boolean"
  },
  "fieldForUserId": {
    "type": "string"
  },
  "timeZoneId": {
    "type": "string",
    "enum": [
      "Africa/Abidjan",
      "Africa/Accra",
      "Africa/Addis_Ababa",
      "Africa/Algiers",
      "Africa/Asmara",
      "Africa/Asmera",
      "Africa/Bamako",
      "Africa/Bangui",
      "Africa/Banjul",
      "Africa/Bissau",
      "Africa/Blantyre",
      "Africa/Brazzaville",
      "Africa/Bujumbura",
      "Africa/Cairo",
      "Africa/Casablanca",
      "Africa/Ceuta",
      "Africa/Conakry",
      "Africa/Dakar",
      "Africa/Dar_es_Salaam",
      "Africa/Djibouti",
      "Africa/Douala",
      "Africa/El_Aaiun",
      "Africa/Freetown",
      "Africa/Gaborone",
      "Africa/Harare",
      "Africa/Johannesburg",
      "Africa/Juba",
      "Africa/Kampala",
      "Africa/Khartoum",
      "Africa/Kigali",
      "Africa/Kinshasa",
      "Africa/Lagos",
      "Africa/Libreville",
      "Africa/Lome",
```



```
"Africa/Luanda",  
"Africa/Lubumbashi",  
"Africa/Lusaka",  
"Africa/Malabo",  
"Africa/Maputo",  
"Africa/Maseru",  
"Africa/Mbabane",  
"Africa/Mogadishu",  
"Africa/Monrovia",  
"Africa/Nairobi",  
"Africa/Ndjamena",  
"Africa/Niamey",  
"Africa/Nouakchott",  
"Africa/Ouagadougou",  
"Africa/Porto-Novo",  
"Africa/Sao_Tome",  
"Africa/Timbuktu",  
"Africa/Tripoli",  
"Africa/Tunis",  
"Africa/Windhoek",  
"America/Adak",  
"America/Anchorage",  
"America/Anguilla",  
"America/Antigua",  
"America/Araguaina",  
"America/Argentina/Buenos_Aires",  
"America/Argentina/Catamarca",  
"America/Argentina/ComodRivadavia",  
"America/Argentina/Cordoba",  
"America/Argentina/Jujuy",  
"America/Argentina/La_Rioja",  
"America/Argentina/Mendoza",  
"America/Argentina/Rio_Gallegos",  
"America/Argentina/Salta",  
"America/Argentina/San_Juan",  
"America/Argentina/San_Luis",  
"America/Argentina/Tucuman",  
"America/Argentina/Ushuaia",  
"America/Aruba",  
"America/Asuncion",  
"America/Atikokan",  
"America/Atka",  
"America/Bahia",  
"America/Bahia_Banderas",
```

```
"America/Barbados",  
"America/Belem",  
"America/Belize",  
"America/Blanc-Sablon",  
"America/Boa_Vista",  
"America/Bogota",  
"America/Boise",  
"America/Buenos_Aires",  
"America/Cambridge_Bay",  
"America/Campo_Grande",  
"America/Cancun",  
"America/Caracas",  
"America/Catamarca",  
"America/Cayenne",  
"America/Cayman",  
"America/Chicago",  
"America/Chihuahua",  
"America/Ciudad_Juarez",  
"America/Coral_Harbour",  
"America/Cordoba",  
"America/Costa_Rica",  
"America/Creston",  
"America/Cuiaba",  
"America/Curacao",  
"America/Danmarkshavn",  
"America/Dawson",  
"America/Dawson_Creek",  
"America/Denver",  
"America/Detroit",  
"America/Dominica",  
"America/Edmonton",  
"America/Eirunepe",  
"America/El_Salvador",  
"America/Ensenada",  
"America/Fort_Nelson",  
"America/Fort_Wayne",  
"America/Fortaleza",  
"America/Glace_Bay",  
"America/Godthab",  
"America/Goose_Bay",  
"America/Grand_Turk",  
"America/Grenada",  
"America/Guadeloupe",  
"America/Guatemala",
```

```
"America/Guayaquil",
"America/Guyana",
"America/Halifax",
"America/Havana",
"America/Hermosillo",
"America/Indiana/Indianapolis",
"America/Indiana/Knox",
"America/Indiana/Marengo",
"America/Indiana/Petersburg",
"America/Indiana/Tell_City",
"America/Indiana/Vevay",
"America/Indiana/Vincennes",
"America/Indiana/Winamac",
"America/Indianapolis",
"America/Inuvik",
"America/Iqaluit",
"America/Jamaica",
"America/Jujuy",
"America/Juneau",
"America/Kentucky/Louisville",
"America/Kentucky/Monticello",
"America/Knox_IN",
"America/Kralendijk",
"America/La_Paz",
"America/Lima",
"America/Los_Angeles",
"America/Louisville",
"America/Lower_Princes",
"America/Maceio",
"America/Managua",
"America/Manaus",
"America/Marigot",
"America/Martinique",
"America/Matamoros",
"America/Mazatlan",
"America/Mendoza",
"America/Menominee",
"America/Merida",
"America/Metlakatla",
"America/Mexico_City",
"America/Miquelon",
"America/Moncton",
"America/Monterrey",
"America/Montevideo",
```

```
"America/Montreal",
"America/Montserrat",
"America/Nassau",
"America/New_York",
"America/Nipigon",
"America/Nome",
"America/Noronha",
"America/North_Dakota/Beulah",
"America/North_Dakota/Center",
"America/North_Dakota/New_Salem",
"America/Nuuk",
"America/Ojinaga",
"America/Panama",
"America/Pangnirtung",
"America/Paramaribo",
"America/Phoenix",
"America/Port-au-Prince",
"America/Port_of_Spain",
"America/Porto_Acre",
"America/Porto_Velho",
"America/Puerto_Rico",
"America/Punta_Arenas",
"America/Rainy_River",
"America/Rankin_Inlet",
"America/Recife",
"America/Regina",
"America/Resolute",
"America/Rio_Branco",
"America/Rosario",
"America/Santa_Isabel",
"America/Santarem",
"America/Santiago",
"America/Santo_Domingo",
"America/Sao_Paulo",
"America/Scoresbysund",
"America/Shiprock",
"America/Sitka",
"America/St_Barthelemy",
"America/St_Johns",
"America/St_Kitts",
"America/St_Lucia",
"America/St_Thomas",
"America/St_Vincent",
"America/Swift_Current",
```

```
"America/Tegucigalpa",
"America/Thule",
"America/Thunder_Bay",
"America/Tijuana",
"America/Toronto",
"America/Tortola",
"America/Vancouver",
"America/Virgin",
"America/Whitehorse",
"America/Winnipeg",
"America/Yakutat",
"America/Yellowknife",
"Antarctica/Casey",
"Antarctica/Davis",
"Antarctica/DumontDURville",
"Antarctica/Macquarie",
"Antarctica/Mawson",
"Antarctica/McMurdo",
"Antarctica/Palmer",
"Antarctica/Rothera",
"Antarctica/South_Pole",
"Antarctica/Syowa",
"Antarctica/Troll",
"Antarctica/Vostok",
"Arctic/Longyearbyen",
"Asia/Aden",
"Asia/Almaty",
"Asia/Amman",
"Asia/Anadyr",
"Asia/Aqtau",
"Asia/Aqtobe",
"Asia/Ashgabat",
"Asia/Ashkhabad",
"Asia/Atyrau",
"Asia/Baghdad",
"Asia/Bahrain",
"Asia/Baku",
"Asia/Bangkok",
"Asia/Barnaul",
"Asia/Beirut",
"Asia/Bishkek",
"Asia/Brunei",
"Asia/Calcutta",
"Asia/Chita",
```

```
"Asia/Choibalsan",  
"Asia/Chongqing",  
"Asia/Chungking",  
"Asia/Colombo",  
"Asia/Dacca",  
"Asia/Damascus",  
"Asia/Dhaka",  
"Asia/Dili",  
"Asia/Dubai",  
"Asia/Dushanbe",  
"Asia/Famagusta",  
"Asia/Gaza",  
"Asia/Harbin",  
"Asia/Hebron",  
"Asia/Ho_Chi_Minh",  
"Asia/Hong_Kong",  
"Asia/Hovd",  
"Asia/Irkutsk",  
"Asia/Istanbul",  
"Asia/Jakarta",  
"Asia/Jayapura",  
"Asia/Jerusalem",  
"Asia/Kabul",  
"Asia/Kamchatka",  
"Asia/Karachi",  
"Asia/Kashgar",  
"Asia/Kathmandu",  
"Asia/Katmandu",  
"Asia/Khandyga",  
"Asia/Kolkata",  
"Asia/Krasnoyarsk",  
"Asia/Kuala_Lumpur",  
"Asia/Kuching",  
"Asia/Kuwait",  
"Asia/Macao",  
"Asia/Macau",  
"Asia/Magadan",  
"Asia/Makassar",  
"Asia/Manila",  
"Asia/Muscat",  
"Asia/Nicosia",  
"Asia/Novokuznetsk",  
"Asia/Novosibirsk",  
"Asia/Omsk",
```

```
"Asia/Oral",
"Asia/Phnom_Penh",
"Asia/Pontianak",
"Asia/Pyongyang",
"Asia/Qatar",
"Asia/Qostanay",
"Asia/Qyzylorda",
"Asia/Rangoon",
"Asia/Riyadh",
"Asia/Saigon",
"Asia/Sakhalin",
"Asia/Samarkand",
"Asia/Seoul",
"Asia/Shanghai",
"Asia/Singapore",
"Asia/Srednekolymsk",
"Asia/Taipei",
"Asia/Tashkent",
"Asia/Tbilisi",
"Asia/Tehran",
"Asia/Tel_Aviv",
"Asia/Thimbu",
"Asia/Thimphu",
"Asia/Tokyo",
"Asia/Tomsk",
"Asia/Ujung_Pandang",
"Asia/Ulaanbaatar",
"Asia/Ulan_Bator",
"Asia/Urumqi",
"Asia/Ust-Nera",
"Asia/Vientiane",
"Asia/Vladivostok",
"Asia/Yakutsk",
"Asia/Yangon",
"Asia/Yekaterinburg",
"Asia/Yerevan",
"Atlantic/Azores",
"Atlantic/Bermuda",
"Atlantic/Canary",
"Atlantic/Cape_Verde",
"Atlantic/Faeroe",
"Atlantic/Faroe",
"Atlantic/Jan_Mayen",
"Atlantic/Madeira",
```

```
"Atlantic/Reykjavik",  
"Atlantic/South_Georgia",  
"Atlantic/St_Helena",  
"Atlantic/Stanley",  
"Australia/ACT",  
"Australia/Adelaide",  
"Australia/Brisbane",  
"Australia/Broken_Hill",  
"Australia/Canberra",  
"Australia/Currie",  
"Australia/Darwin",  
"Australia/Eucla",  
"Australia/Hobart",  
"Australia/LHI",  
"Australia/Lindeman",  
"Australia/Lord_Howe",  
"Australia/Melbourne",  
"Australia/NSW",  
"Australia/North",  
"Australia/Perth",  
"Australia/Queensland",  
"Australia/South",  
"Australia/Sydney",  
"Australia/Tasmania",  
"Australia/Victoria",  
"Australia/West",  
"Australia/Yancowinna",  
"Brazil/Acre",  
"Brazil/DeNoronha",  
"Brazil/East",  
"Brazil/West",  
"CET",  
"CST6CDT",  
"Canada/Atlantic",  
"Canada/Central",  
"Canada/Eastern",  
"Canada/Mountain",  
"Canada/Newfoundland",  
"Canada/Pacific",  
"Canada/Saskatchewan",  
"Canada/Yukon",  
"Chile/Continental",  
"Chile/EasterIsland",  
"Cuba",
```



```
"EET",  
"EST5EDT",  
"Egypt",  
"Eire",  
"Etc/GMT",  
"Etc/GMT+0",  
"Etc/GMT+1",  
"Etc/GMT+10",  
"Etc/GMT+11",  
"Etc/GMT+12",  
"Etc/GMT+2",  
"Etc/GMT+3",  
"Etc/GMT+4",  
"Etc/GMT+5",  
"Etc/GMT+6",  
"Etc/GMT+7",  
"Etc/GMT+8",  
"Etc/GMT+9",  
"Etc/GMT-0",  
"Etc/GMT-1",  
"Etc/GMT-10",  
"Etc/GMT-11",  
"Etc/GMT-12",  
"Etc/GMT-13",  
"Etc/GMT-14",  
"Etc/GMT-2",  
"Etc/GMT-3",  
"Etc/GMT-4",  
"Etc/GMT-5",  
"Etc/GMT-6",  
"Etc/GMT-7",  
"Etc/GMT-8",  
"Etc/GMT-9",  
"Etc/GMT0",  
"Etc/Greenwich",  
"Etc/UCT",  
"Etc/UTC",  
"Etc/Universal",  
"Etc/Zulu",  
"Europe/Amsterdam",  
"Europe/Andorra",  
"Europe/Astrakhan",  
"Europe/Athens",  
"Europe/Belfast",
```

```
"Europe/Belgrade",  
"Europe/Berlin",  
"Europe/Bratislava",  
"Europe/Brussels",  
"Europe/Bucharest",  
"Europe/Budapest",  
"Europe/Busingen",  
"Europe/Chisinau",  
"Europe/Copenhagen",  
"Europe/Dublin",  
"Europe/Gibraltar",  
"Europe/Guernsey",  
"Europe/Helsinki",  
"Europe/Isle_of_Man",  
"Europe/Istanbul",  
"Europe/Jersey",  
"Europe/Kaliningrad",  
"Europe/Kiev",  
"Europe/Kirov",  
"Europe/Kyiv",  
"Europe/Lisbon",  
"Europe/Ljubljana",  
"Europe/London",  
"Europe/Luxembourg",  
"Europe/Madrid",  
"Europe/Malta",  
"Europe/Mariehamn",  
"Europe/Minsk",  
"Europe/Monaco",  
"Europe/Moscow",  
"Europe/Nicosia",  
"Europe/Oslo",  
"Europe/Paris",  
"Europe/Podgorica",  
"Europe/Prague",  
"Europe/Riga",  
"Europe/Rome",  
"Europe/Samara",  
"Europe/San_Marino",  
"Europe/Sarajevo",  
"Europe/Saratov",  
"Europe/Simferopol",  
"Europe/Skopje",  
"Europe/Sofia",
```

```
"Europe/Stockholm",  
"Europe/Tallinn",  
"Europe/Tirane",  
"Europe/Tiraspol",  
"Europe/Ulyanovsk",  
"Europe/Uzhgorod",  
"Europe/Vaduz",  
"Europe/Vatican",  
"Europe/Vienna",  
"Europe/Vilnius",  
"Europe/Volgograd",  
"Europe/Warsaw",  
"Europe/Zagreb",  
"Europe/Zaporozhye",  
"Europe/Zurich",  
"GB",  
"GB-Eire",  
"GMT",  
"GMT0",  
"Greenwich",  
"Hongkong",  
"Iceland",  
"Indian/Antananarivo",  
"Indian/Chagos",  
"Indian/Christmas",  
"Indian/Cocos",  
"Indian/Comoro",  
"Indian/Kerguelen",  
"Indian/Mahe",  
"Indian/Maldives",  
"Indian/Mauritius",  
"Indian/Mayotte",  
"Indian/Reunion",  
"Iran",  
"Israel",  
"Jamaica",  
"Japan",  
"Kwajalein",  
"Libya",  
"MET",  
"MST7MDT",  
"Mexico/BajaNorte",  
"Mexico/BajaSur",  
"Mexico/General",
```

```
"NZ",  
"NZ-CHAT",  
"Navajo",  
"PRC",  
"PST8PDT",  
"Pacific/Apia",  
"Pacific/Auckland",  
"Pacific/Bougainville",  
"Pacific/Chatham",  
"Pacific/Chuuk",  
"Pacific/Easter",  
"Pacific/Efate",  
"Pacific/Enderbury",  
"Pacific/Fakaofu",  
"Pacific/Fiji",  
"Pacific/Funafuti",  
"Pacific/Galapagos",  
"Pacific/Gambier",  
"Pacific/Guadalcanal",  
"Pacific/Guam",  
"Pacific/Honolulu",  
"Pacific/Johnston",  
"Pacific/Kanton",  
"Pacific/Kiritimati",  
"Pacific/Kosrae",  
"Pacific/Kwajalein",  
"Pacific/Majuro",  
"Pacific/Marquesas",  
"Pacific/Midway",  
"Pacific/Nauru",  
"Pacific/Niue",  
"Pacific/Norfolk",  
"Pacific/Noumea",  
"Pacific/Pago_Pago",  
"Pacific/Palau",  
"Pacific/Pitcairn",  
"Pacific/Pohnpei",  
"Pacific/Ponape",  
"Pacific/Port_Moresby",  
"Pacific/Rarotonga",  
"Pacific/Saipan",  
"Pacific/Samoa",  
"Pacific/Tahiti",  
"Pacific/Tarawa",
```

```
"Pacific/Tongatapu",  
"Pacific/Truk",  
"Pacific/Wake",  
"Pacific/Wallis",  
"Pacific/Yap",  
"Poland",  
"Portugal",  
"ROK",  
"Singapore",  
"SystemV/AST4",  
"SystemV/AST4ADT",  
"SystemV/CST6",  
"SystemV/CST6CDT",  
"SystemV/EST5",  
"SystemV/EST5EDT",  
"SystemV/HST10",  
"SystemV/MST7",  
"SystemV/MST7MDT",  
"SystemV/PST8",  
"SystemV/PST8PDT",  
"SystemV/YST9",  
"SystemV/YST9YDT",  
"Turkey",  
"UCT",  
"US/Alaska",  
"US/Aleutian",  
"US/Arizona",  
"US/Central",  
"US/East-Indiana",  
"US/Eastern",  
"US/Hawaii",  
"US/Indiana-Starke",  
"US/Michigan",  
"US/Mountain",  
"US/Pacific",  
"US/Samoa",  
"UTC",  
"Universal",  
"W-SU",  
"WET",  
"Zulu",  
"EST",  
"HST",  
"MST",
```

```
    "ACT",
    "AET",
    "AGT",
    "ART",
    "AST",
    "BET",
    "BST",
    "CAT",
    "CNT",
    "CST",
    "CTT",
    "EAT",
    "ECT",
    "IET",
    "IST",
    "JST",
    "MIT",
    "NET",
    "NST",
    "PLT",
    "PNT",
    "PRT",
    "PST",
    "SST",
    "VST"
  ]
},
"pageRootPaths": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"assetRootPaths": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"crawlAssets": {
  "type": "boolean"
},
"crawlPages": {
  "type": "boolean"
```

```
    },
    "pagePathInclusionPatterns": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    "pagePathExclusionPatterns": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    "pageNameInclusionPatterns": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    "pageNameExclusionPatterns": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    "assetPathInclusionPatterns": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    "assetPathExclusionPatterns": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    "assetTypeInclusionPatterns": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    "assetTypeExclusionPatterns": {
```

```
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "assetNameInclusionPatterns": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "assetNameExclusionPatterns": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "pageComponents": {
    "type": "array",
    "items": {
      "type": "object"
    }
  },
  "contentFragmentVariations": {
    "type": "array",
    "items": {
      "type": "object"
    }
  },
  "cugExemptedPrincipals": {
    "type": "array",
    "items": {
      "type": "string"
    }
  }
},
"required": [],
"type": {
  "type": "string",
  "pattern": "AEM"
},
"enableIdentityCrawler": {
  "type": "boolean"
```



```

    },
    "syncMode": {
      "type": "string",
      "enum": [
        "FORCED_FULL_CRAWL",
        "FULL_CRAWL",
        "CHANGE_LOG"
      ]
    },
    "secretArn": {
      "type": "string",
      "minLength": 20,
      "maxLength": 2048
    }
  },
  "version": {
    "type": "string",
    "anyOf": [
      {
        "pattern": "1.0.0"
      }
    ]
  },
  "required": [
    "connectionConfiguration",
    "repositoryConfigurations",
    "syncMode",
    "additionalProperties",
    "secretArn",
    "type"
  ]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.
repositoryEndpointMetadata	The endpoint information for the data source.

Configuration	Description
<code>aemUrl</code>	The Adobe Experience Manager host URL. For example, if you use AEM On-Premise, you include the hostname and port: <i>http://hostname:port</i> . Or, if you use AEM as a Cloud Service, you can use the author URL: <i>https://author-xxxxxx-xxxxxxx.adobecloud.com</i> .
<code>authType</code>	The type of authentication you use, whether Basic or OAuth2.
<code>deploymentType</code>	The type of Adobe Experience Manager that you use, either CLOUD or ON-PREMISE .
<code>repositoryConfigurations</code>	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings.
<ul style="list-style-type: none"> <code>page</code> <code>asset</code> 	A list of objects that map the attributes or field names of your Adobe Experience Manager pages and assets to Amazon Q index field names.
<code>additionalProperties</code>	Additional configuration options for your content in your data source.
<code>isCrawlAcl</code>	Specify <code>true</code> to crawl access control information from documents.
<code>fieldForUserId</code>	Specify field to use for <code>UserId</code> for ACL crawling.

Configuration	Description
timeZoneId	<p>If you use AEM On-Premise and the time zone of your server is different than the time zone of the Amazon Q AEM connector or index, you can specify the server time zone to align with the AEM connector or index.</p> <p>The default time zone for AEM On-Premise is the time zone of the Amazon Q AEM connector or index. The default time zone for AEM as a Cloud Service is Greenwich Mean Time.</p>
<ul style="list-style-type: none"> • pageRootPaths • assetRootPaths 	<p>A list of root paths for pages and assets. For example, the root path for a page could be <i>/content/sub</i> and the root path for an asset could be <i>/content/sub/asset1</i>.</p>
crawlAssets	Specify true to crawl assets.
crawlPages	Specify true to crawl pages.
<ul style="list-style-type: none"> • pagePathInclusionPatterns • pageNameInclusionPatterns • assetPathInclusionPatterns • assetTypeInclusionPatterns • assetNameInclusionPatterns 	<p>A list of regular expression patterns to include certain pages and assets in your Adobe Experience Manager data source. Pages and assets that match the patterns are included in the index. Pages and assets that don't match the patterns are excluded from the index. If a page or asset matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence, and the content isn't included in the index.</p>

Configuration	Description
<ul style="list-style-type: none"> • <code>pagePathExclusionPatterns</code> • <code>pageNameExclusionPatterns</code> • <code>assetPathExclusionPatterns</code> • <code>assetTypeInclusionPatterns</code> • <code>assetNameInclusionPatterns</code> 	<p>A list of regular expression patterns to exclude certain pages and assets in your Adobe Experience Manager data source. Pages and assets that match the patterns are excluded from the index. Pages and assets that don't match the patterns are included in the index. If a page or asset matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence, and the content isn't included in the index.</p>
<code>pageComponents</code>	<p>A list of names for the specific page components that you want to index.</p>
<code>contentFragmentVariations</code>	<p>A list of names for the specific saved variations of Adobe Experience Manager Content Fragments that you want to index.</p>
<code>type</code>	<p>The type of data source. Specify AEM as your data source type.</p>
<code>enableIdentityCrawler</code>	<p>Specify <code>true</code> to use the Amazon Q identity crawler to sync identity/principal information on users and groups with access to specific documents. See Identity crawler for more information.</p>

Configuration	Description
syncMode	<p>Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose between the following options:</p> <ul style="list-style-type: none">• Use <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index.• Use <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index.• Use <code>CHANGE_LOG</code> to incrementally crawl only new and modified content each time your data source syncs with your index.

Configuration	Description
secretArn	<p>The Amazon Resource Name (ARN) of an AWS Secrets Manager secret that contains the key-value pairs required to connect to your Adobe Experience Manager. The secret must contain a JSON structure with the following keys:</p> <p>If using basic authentication for either AEM On-Premise or Cloud:</p> <pre data-bbox="829 617 1507 974">{ "aemUrl": "<i>Adobe Experience Manager On-Premise host URL</i> ", "username": "<i>user name with admin permissions</i> ", "password": "<i>password with admin permissions</i> " }</pre> <p>If using OAuth 2.0 authentication for AEM On-Premise:</p> <pre data-bbox="829 1129 1507 1446">{ "aemUrl": "<i>Adobe Experience Manager host URL</i>", "clientId": "<i>client ID</i>", "clientSecret": "<i>client secret</i>", "privateKey": "<i>private key</i>" }</pre> <p>If using OAuth 2.0 authentication for AEM as a Cloud Service:</p> <pre data-bbox="829 1602 1507 1820">{ "clientId": "<i>client ID</i>", "clientSecret": "<i>client secret</i>", "privateKey": "<i>private key</i>", "orgId": "<i>organization ID</i> ", }</pre>

Configuration	Description
	<pre data-bbox="828 205 1510 420"> "technicalAccountId": " <i>technical account ID</i>", "imsHost": " <i>Adobe Identity Management System (IMS) host</i> " } </pre>
version	The version of this template that's currently supported.

How Amazon Q Business connector crawls AEM (Cloud) ACLs

When you connect an AEM (Cloud) data source to Amazon Q Business, Amazon Q Business crawls ACL information attached to a document (user and group information) from your AEM (Cloud) instance. If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

The group and user IDs are mapped as follows:

- `_group_ids` – Group IDs exist in Adobe Experience Manager content where there are set access permissions. They're mapped from the names of the groups in AEM.
- `_user_id` – User IDs exist in Adobe Experience Manager content where there are set access permissions. They're mapped from the user emails as the IDs in AEM.

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business AEM (Cloud) data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q Business enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a attribute mapping already available, or if you want to map additional document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q Adobe Experience Manager (AEM) connector supports the following entities and the associated reserved and custom attributes.

Supported entities and field mappings

- [Pages](#)
- [Assets](#)

Pages

Amazon Q supports crawling [AEM Pages](#) and offers the following page field mappings.

Adobe Experience Manager (AEM) field name	Index field name	Description	Data type
aem_page_source_uri	_source_uri	Default	String
aem_page_createdBy	_authors	Default	String list
aem_page_template	aem_page_template	Custom	String
aem_entity_type	_category	Default	String
aem_page_createdAt	_created_at	Default	Date
aem_page_lastModified	_last_updated_at	Default	Date
aem_page_lastReplicatedBy	aem_page_publisher	Custom	String
aem_page_lastReplicatedAt	aem_page_publishedAt	Custom	Date

Assets

Amazon Q supports crawling [AEM Assets](#) and offers the following asset field mappings.

Adobe Experience Manager (AEM) field name	Index field name	Description	Data type
aem_page_source_uri	_source_uri	Default	String
aem_page_createdBy	_authors	Default	String list
aem_entity_type	_category	Default	String
aem_page_createdAt	_created_at	Default	Date

Adobe Experience Manager (AEM) field name	Index field name	Description	Data type
aem_page_lastModified	_last_updated_at	Default	Date
aem_page_lastRepliatedBy	aem_page_publisher	Custom	String
aem_page_lastRepliatedAt	aem_page_published At	Custom	Date

IAM role for Amazon Q AEM (Cloud) connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the `BatchPutDocument` and `BatchDeleteDocument` operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
```

```

    "Sid": "AllowsAmazonQToGetSecret",
    "Effect": "Allow",
    "Action": [
      "secretsmanager:GetSecretValue"
    ],
    "Resource": [
      "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToDecryptSecret",
    "Effect": "Allow",
    "Action": [
      "kms:Decrypt"
    ],
    "Resource": [
      "arn:aws:kms:{{region}}:{{account_id}}:key/[key_id]"
    ],
    "Condition": {
      "StringLike": {
        "kms:ViaService": [
          "secretsmanager.*.amazonaws.com"
        ]
      }
    }
  }
},
{
  "Sid": "AllowsAmazonQToIngestDocuments",
  "Effect": "Allow",
  "Action": [
    "qbusiness:BatchPutDocument",
    "qbusiness:BatchDeleteDocument"
  ],
  "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
},
{
  "Sid": "AllowsAmazonQToIngestPrincipalMapping",
  "Effect": "Allow",
  "Action": [
    "qbusiness:PutGroup",
    "qbusiness:CreateUser",
    "qbusiness>DeleteGroup",
    "qbusiness:UpdateUser",

```

```

    "qbusiness:ListGroupsWith"
  ],
  "Resource": [
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}",
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}",
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}/data-source/*"
  ]
},
{
  "Sid": "AllowsAmazonQToCreateAndDeleteNI",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateNetworkInterface",
    "ec2:DeleteNetworkInterface"
  ],
  "Resource": [
    "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[subnet_ids]",
    "arn:aws:ec2:{{region}}:{{account_id}}:security-group/[security_group]"
  ]
},
{
  "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateNetworkInterface",
    "ec2:DeleteNetworkInterface"
  ],
  "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
  "Condition": {
    "StringLike": {
      "aws:RequestTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
    },
    "ForAllValues:StringEquals": {
      "aws:TagKeys": [
        "AMAZON_Q"
      ]
    }
  }
},
{
  "Sid": "AllowsAmazonQToCreateTags",
  "Effect": "Allow",

```

```

    "Action": [
      "ec2:CreateTags"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringEquals": {
        "ec2:CreateAction": "CreateNetworkInterface"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterfacePermission"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:ResourceTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
    "Effect": "Allow",
    "Action": [
      "ec2:DescribeNetworkInterfaces",
      "ec2:DescribeAvailabilityZones",
      "ec2:DescribeNetworkInterfaceAttribute",
      "ec2:DescribeVpcs",
      "ec2:DescribeRegions",
      "ec2:DescribeNetworkInterfacePermissions",
      "ec2:DescribeSubnets"
    ],
    "Resource": "*"
  }
]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```
{
```

```

"Version": "2012-10-17",
"Statement": [
  {
    "Sid": "AllowsAmazonQServicePrincipal",
    "Effect": "Allow",
    "Principal": {
      "Service": "qbusiness.amazonaws.com"
    },
    "Action": "sts:AssumeRole",
    "Condition": {
      "StringEquals": {
        "aws:SourceAccount": "{{source_account}}"
      },
      "ArnEquals": {
        "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
      }
    }
  }
]
}

```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Known limitations for the Amazon Q Business AEM (Cloud) connector

The Amazon Q Business AEM (Cloud) connector has the following known limitations:

- Deleted site pages can't be tracked when you use **Change log sync** or **Sync only new, modified, or deleted document sync**.

Troubleshooting your Amazon Q Business AEM (Cloud) connector

The following table provides information about error codes you may see for the Adobe Experience Manager (AEM) connector and suggested troubleshooting actions.

Error code	Error message	Suggested resolution
AEM-5001	Error while getting Administrators group.	Check whether provided username and password are correct or not. Also

Error code	Error message	Suggested resolution
	<p>Below are the possible reasons for this error:</p> <p>Provided AEM host URL might be wrong.</p> <p>Provided username and password are invalid or user is non-admin user.</p>	<p>ensure that the provided user is either admin or belongs to administrators' group.</p>
AEM-5002	Error while generating OAuth2 access token.	Provide valid OAuth2 credentials.
AEM-5103	Null/empty AEM host URL.	AEM host URL should not be null or empty.
AEM-5104	<p>Error while parsing https response. Below are the possible reasons for this error.</p> <ol style="list-style-type: none"> 1. Provided AEM host URL might be wrong, please cross-check the AEM host URL. 2. AEM server is down or not reachable. 	Provide a valid host URL, or try again later.
AEM-5105	Provided authType is incorrect.	Auth type should be Basic or OAuth2.
AEM-5106	Null/empty AEM username.	Username should not be null or empty value.
AEM-5107	Null/empty AEM password.	Password should not be null or empty value.

Error code	Error message	Suggested resolution
AEM-5108	Null/empty client id.	Client Id should not be null or empty value.
AEM-5109	Null/empty client secret.	Client Secret should not be null or empty value.
AEM-5110	Null/empty private key.	Private key should not be null or empty value.
AEM-5111	Null/empty Page Index field name.	Page index field should not be null or empty value
AEM-5112	Null/empty Page data source field name.	Page data source field should not be null or empty value.
AEM-5113	Null/empty asset Index field name.	Asset index field should not be null or empty value.
AEM-5114	Null/empty asset data source field name.	Asset data source field should not be null or empty value.
AEM-5115	Null/empty crawl type.	crawl Type value should be FULL_CRAWL/CHANG_LOG type.
AEM-5116	Invalid AEM host URL format.	Check whether provided AEM URL is in correct format or not e.g. http<s>://<aem-host>:<port>
AEM-5117	Page root paths are incorrect.	Page root paths must be a list of strings.
AEM-5118	Asset root paths are incorrect.	Asset root paths must be a list of strings.
AEM-5119	Page path inclusion or exclusion patterns are incorrect.	Page name inclusion patterns/ Exclusion must be a list of strings.

Error code	Error message	Suggested resolution
AEM-5120	Asset path inclusion or exclusion patterns are incorrect.	Asset name inclusion patterns/ Exclusion must be a list of strings.
AEM-5121	Provided deploymenttype is incorrect.	Deployment type should be either CLOUD or ON_PREMISE.
AEM-5122	Provided orgId is incorrect.	OrgId should not be null or empty value.
AEM-5123	Provided technical Account Id is incorrect.	Technical Account Id should not be null or empty value.
AEM-5124	Provided imsHost is incorrect.	IMS Host should not be null or empty value.
AEM-5125	Null/Empty deployment type.	Deployment type should be either CLOUD or ON_PREMISE.
AEM-5126	Invalid Timezone Id.	Provide a valid timezone id.
AEM-5127	Null/empty asset Index field type.	Asset index field should not be null or empty value.
AEM-5128	Null/empty page Index field type.	Page index field should not be null or empty value.
AEM-5129	DataSourceFieldName doesn't match with IndexFieldType.	Provide a valid asset indexFieldType for the provided asset dataSourceFieldName. Or, provide a valid page indexFieldType for the provided page dataSourceFieldName.
AEM-5130	Protocol used by provided AEM URL is not supported by AEM connector.	Only https protocol is supported by AEM connector. Provide an AEM URL based on https protocol.

Error code	Error message	Suggested resolution
AEM-5131	AEM password is too large.	Password should not be greater than 40 characters.
AEM-5132	AEM client ID is too large.	Client ID should not be greater than 40 characters.
AEM-5133	AEM client secret is too large.	Client secret should not be greater than 40 characters.
AEM-5134	AEM private key is too large.	Private key should not be greater than 2048 characters.
AEM-5135	AEM client ID contains invalid characters.	Client ID should not contain unprintable characters.
AEM-5136	AEM client secret contains invalid characters.	Client secret should not contain unprintable characters.
AEM-5137	AEM private key contains invalid characters.	Private key should not contain unprintable characters.
AEM-5138	AEM IMS host is too large.	IMS host should not be greater than 100 characters.
AEM-5139	AEM technical account ID is too large.	Technical account id should not be greater than 100 characters.
AEM-5140	AEM org ID is too large.	Org id should not be greater than 100 characters.
AEM-5141	Page name inclusion or exclusion patterns are incorrect.	Page name inclusion patterns/ Exclusion must be a list of strings.

Error code	Error message	Suggested resolution
AEM-5142	Asset name inclusion or exclusion patterns are incorrect.	Asset name inclusion patterns/ Exclusion must be a list of strings.
AEM-5143	Asset type inclusion or exclusion patterns are incorrect.	Asset type inclusion patterns/Exclusion must be a list of strings.
AEM-5144	Invalid page root path. Please provide valid page root path.	Page path should start with /content.
AEM-5145	Invalid asset root path. Please provide valid asset root path.	Asset path should start with /content/ dam.
AEM-5146	AEM page root paths list size is too large.	Page root paths list size should not be greater than 1000.
AEM-5147	AEM asset root paths list size is too large.	Asset root paths list size should not be greater than 1000.
AEM-5148	Asset root paths list size should not be greater than 1000.	Asset path exclusion patterns list size should not be greater than 1000.
AEM-5149	AEM asset path inclusion pattern list size is too large.	Asset path inclusion patterns list size should not be greater than 1000.
AEM-5150	AEM asset name inclusion pattern list size is too large.	Asset name inclusion patterns list size should not be greater than 1000.
AEM-5151	AEM asset name exclusion pattern list size is too large.	Asset name exclusion patterns list size should not be greater than 1000.

Error code	Error message	Suggested resolution
AEM-5152	AEM asset type exclusion pattern list size is too large.	Asset type exclusion patterns list size should not be greater than 1000.
AEM-5153	AEM asset type inclusion pattern list size is too large.	Asset type inclusion patterns list size should not be greater than 1000.
AEM-5154	AEM page name inclusion pattern list size is too large.	Page name inclusion patterns list size should not be greater than 1000.
AEM-5155	AEM page name Exclusion pattern list size is too large.	Page name exclusion patterns list size should not be greater than 1000.
AEM-5156	AEM page path Exclusion pattern list size is too large.	Page path exclusion patterns list size should not be greater than 1000.
AEM-5157	AEM page path inclusion pattern list size is too large.	Page path inclusion patterns list size should not be greater than 1000.
AEM-5158	AEM page components list size is too large.	Page components list size should not be greater than 1000.
AEM-5159	AEM content fragment variations list size is too large.	Content fragment variations list size should not be greater than 1000.
AEM-5160	AEM host URL characters length is too large.	AEM host URL characters length should not be greater than 2048 characters.
AEM-5161	Some of the page root paths exceed the character limit.	Page root path characters length should not be greater than 1000.

Error code	Error message	Suggested resolution
AEM-5162	Some of the asset root paths exceed the character limit.	Asset root Path characters length should not be greater than 1000 .
AEM-5163	Some of the asset path exclusion objects exceed the character limit.	Asset path exclusion characters length should not be greater than 1000.
AEM-5164	Some of the asset path inclusion objects exceed the character limit.	Asset path inclusion characters length should not be greater than 1000.
AEM-5165	Some of the asset name inclusion objects exceed the character limit.	Asset name inclusion characters length should not be greater than 1000.
AEM-5166	Some of the asset name exclusion objects exceed the character limit.	Asset name exclusion characters length should not be greater than 1000.
AEM-5167	Some of the asset type exclusion objects exceed the character limit.	Asset type exclusion characters length should not be greater than 1000.
AEM-5168	Some of the asset type inclusion objects exceed the character limit.	Asset type inclusion characters length should not be greater than 1000.
AEM-5169	Some of the page name inclusion objects exceed the character limit.	Page name inclusion characters length should not be greater than 1000.
AEM-5170	Some of the page name exclusion objects exceed the character limit.	Page name exclusion characters length should not be greater than 1000.

Error code	Error message	Suggested resolution
AEM-5171	Some of the page path exclusion objects exceed the character limit.	Page path exclusion characters length should not be greater than 1000.
AEM-5172	Some of the page path inclusion objects exceed the character limit.	Page path inclusion characters length should not be greater than 1000.
AEM-5300	Error in serializing change log token.	Retry sync.
AEM-5301	Error in de-serializing change log token.	Retry sync.
AEM-5401	Error occurred while getting AEM groups.	Retry sync.
AEM-5501	Could not connect to host.	
AEM-5502	AEM URL SSRF check failed.	Make sure AEM host URL is not a multicast/local/link-local/loopback address.
AEM-5503	AEM host not found.	Check whether AEM host is up and reachable.
AEM-5504	Error occurred while executing HTTP request against given AEM URL.	Check whether AEM host is up and reachable.
AEM-5505	AEM malformed URL error.	Provide valid AEM url.
AEM-5506	AEM VPC Configuration check failed.	Site local address is restricted.

Error code	Error message	Suggested resolution
AEM-5507	Error in creating document attribute.	Only String, String List, Date and Long formats are supported for field mappings.
AEM-5200	Error occurred while getting pages from AEM for Full Crawl.	Check whether AEM server is up and responding to API requests.
AEM-5506	AEM VPC Configuration check failed.	Site local address is restricted.
AEM-5201	Error occurred while getting assets from AEM for Full Crawl.	Check whether AEM server is up and responding to API requests.
AEM-5303	Error occurred while getting pages from AEM for Change Log.	Check whether AEM server is up and responding to API requests.
AEM-5304	Error occurred while getting assets from AEM for Change Log.	Check whether AEM server is up and responding to API requests.

Connecting AEM (Server) to Amazon Q Business

Adobe Experience Manager (AEM) is a content management system (CMS) that's used for creating website or mobile app content. You can connect AEM (Server) instance to Amazon Q Business—using either the AWS Management Console or the [CreateDataSource](#) API—and create an Amazon Q web experience.

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).

- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [AEM \(Server\) connector overview](#)
- [Prerequisites for connecting Amazon Q Business to AEM \(Server\)](#)
- [Connecting Amazon Q Business to AEM \(Server\) using the console](#)
- [Connecting Amazon Q Business to AEM \(Server\) using APIs](#)
- [How Amazon Q Business connector crawls AEM \(Server\) ACLs](#)
- [Amazon Q Business AEM \(Server\) data source connector field mappings](#)
- [IAM role for Amazon Q Business AEM \(Server\) connector](#)
- [Known limitations for the Amazon Q Business AEM \(Server\) connector](#)
- [Troubleshooting your Amazon Q Business AEM \(Server\) connector](#)

AEM (Server) connector overview

The following table gives an overview of the Amazon Q Business AEM (Server) connector and its supported features.

Category	Feature	Support
Security	Authentication type	Basic, OAuth 2.0
	Authentication credentials	<p>Basic</p> <ul style="list-style-type: none"> • AEM (Server) host URL • Username of AEM user • Password of AEM user <p>OAuth 2.0</p> <ul style="list-style-type: none"> • AEM (Server) host URL • Client ID • Client secret

Category	Feature	Support
		<ul style="list-style-type: none"> Private key <div style="border: 1px solid #f08080; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p>⚠ Important Admin privileges required.</p> </div>
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Identity crawling	Yes
	VPC	Yes
Crawl features	Custom metadata	Yes
	Entities	Yes. The following entities are supported: <ul style="list-style-type: none"> Pages Assets
	Field mappings	Yes. Supports both default and custom field mappings. For more information, see Field mappings .
	Filters	Yes. The following filters are supported: <ul style="list-style-type: none"> Include/exclude by asset name Include/exclude by asset type Include/exclude by asset path Include/exclude by page name Include/exclude by page path
	Sync mode	Supports full and incremental sync.
	File types	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to AEM (Server)

Before you begin, make sure that you have completed the following prerequisites.

In AEM, make sure you have:

- Access to an account with administrative permissions, or an admin user.
- Copied your AEM host URL.
- Noted your basic authentication credentials of admin username and password.
- (Optional) Added the following OAuth scopes if you're using OAuth 2.0 authentication:
 - **Profile** – Needed to get user and groups related data, like email ID and username.
 - **Replicate** – Needed to get data and metadata from Assets and Pages (not including user data).
- **Optional:** Generated OAuth 2.0 credentials in AEM On-Premise. If you use AEM On-Premise, the credentials include client ID, client secret, and private key. Adobe Granite OAuth 2.0 server implementation (com.adobe.granite.oauth.server) provides the support for OAuth 2.0 server functionalities in AEM.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your AEM (Server) authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Connecting Amazon Q Business to AEM (Server) using the console

The following procedure outlines how to connect Amazon Q Business to AEM (Server) using the AWS Management Console.

Connecting Amazon Q to AEM (Server)

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **AEM (Server)** page, enter the following information:
6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. **Source** – Choose **AEM (Server)** .
 - a. **AEM host URL** – Enter your **AEM host URL**. If you use AEM On-Premise, you include the hostname and port. For example: *https://hostname:port*.
 - b. **SSL certificate location** – Enter the path to the SSL certificate stored in an Amazon S3 bucket. You use this to connect to AEM On-Premise with a secure SSL connection.
8. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

Note

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

9. **Authentication** – Choose between **Basic authentication** and **OAuth 2.0 authentication** and then enter the following information for your **AWS Secrets Manager secret**.
 - a. **Basic authentication** – Enter the name for your secret, your AEM site admin username, and admin password.
 - b. **OAuth 2.0 authentication** – Enter enter a name for the secret, your client ID, client secret, and private key.

10. **Configure VPC and security group – *optional*** – Choose whether you want to use a VPC. If you do, enter the following information:
 - a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
 - b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

11. **Identity crawler** – Choose to activate Amazon Q identity crawler to sync identity information. For more information, see [Identity crawler](#).
12. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

13. In **Sync scope**, enter the following information:
 - a. **Sync content types** – Choose whether to crawl only **Pages** or **Assets**, or both.
 - b. **Additional configuration – *optional*** – Configure the following settings:
 - **Page components** – The specific names of page components. The Page Component is an extensible page component designed to work with the Adobe AEM template editor and allows page header and footer and structure components to be assembled with the template editor.
 - **Content fragment variations** – The specific names of content fragment variations. Content Fragments allow you to design, create, curate and publish page-independent content in Adobe AEM. They allow you to prepare content ready for use in multiple locations and over multiple channels.
 - **Root paths** – The root paths to specific content.
 - **Regex patterns** – The regular expression patterns to include or exclude certain pages and assets.

14. In **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.

- **Full sync** – Sync all content regardless of the previous sync status.
- **New or modified content sync** – Sync only new and modified documents.
- **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.

For more details, see [Sync mode](#).

15. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
16. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
17. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:
 - a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
 - b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

 **Note**

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

18. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

19. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

Note

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to AEM (Server) using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the `configuration` parameter to provide a JSON schema with all other configuration information specific to your data source connector.

For an example of the API request, see [CreateDataSource](#) in the Amazon Q API Reference.

AEM JSON schema

The following is the AEM JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
      "type": "object",
      "properties": {
        "repositoryEndpointMetadata": {
          "type": "object",
          "properties": {
            "aemUrl": {
              "type": "string",
              "pattern": "https:.*"
            },
            "authType": {
              "type": "string",
              "enum": [
                "Basic",
                "OAuth2"
              ]
            }
          }
        }
      }
    }
  }
}
```

```

        "deploymentType": {
            "type": "string",
            "enum": [
                "CLOUD",
                "ON_PREMISE"
            ]
        }
    },
    "required": [
        "aemUrl",
        "authType",
        "deploymentType"
    ]
}
},
"required": [
    "repositoryEndpointMetadata"
]
},
"repositoryConfigurations": {
    "type": "object",
    "properties": {
        "page": {
            "type": "object",
            "properties": {
                "fieldMappings": {
                    "type": "array",
                    "items": [
                        {
                            "type": "object",
                            "properties": {
                                "indexFieldName": {
                                    "type": "string"
                                },
                                "indexFieldType": {
                                    "type": "string",
                                    "enum": [
                                        "STRING",
                                        "STRING_LIST",
                                        "DATE",
                                        "LONG"
                                    ]
                                }
                            }
                        }
                    ]
                },
                "dataSourceFieldName": {

```

```
        "type": "string"
      },
      "dateFieldFormat": {
        "type": "string",
        "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
      }
    },
    "required": [
      "indexFieldName",
      "indexFieldType",
      "dataSourceFieldName"
    ]
  }
]
}
},
"required": [
  "fieldMappings"
]
},
"asset": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": [
        {
          "type": "object",
          "properties": {
            "indexFieldName": {
              "type": "string"
            },
            "indexFieldType": {
              "type": "string",
              "enum": [
                "STRING",
                "STRING_LIST",
                "DATE",
                "LONG"
              ]
            },
            "dataSourceFieldName": {
              "type": "string"
            }
          }
        }
      ]
    }
  }
}
```



```
"Africa/Bujumbura",  
"Africa/Cairo",  
"Africa/Casablanca",  
"Africa/Ceuta",  
"Africa/Conakry",  
"Africa/Dakar",  
"Africa/Dar_es_Salaam",  
"Africa/Djibouti",  
"Africa/Douala",  
"Africa/El_Aaiun",  
"Africa/Freetown",  
"Africa/Gaborone",  
"Africa/Harare",  
"Africa/Johannesburg",  
"Africa/Juba",  
"Africa/Kampala",  
"Africa/Khartoum",  
"Africa/Kigali",  
"Africa/Kinshasa",  
"Africa/Lagos",  
"Africa/Libreville",  
"Africa/Lome",  
"Africa/Luanda",  
"Africa/Lubumbashi",  
"Africa/Lusaka",  
"Africa/Malabo",  
"Africa/Maputo",  
"Africa/Maseru",  
"Africa/Mbabane",  
"Africa/Mogadishu",  
"Africa/Monrovia",  
"Africa/Nairobi",  
"Africa/Ndjamena",  
"Africa/Niamey",  
"Africa/Nouakchott",  
"Africa/Ouagadougou",  
"Africa/Porto-Novo",  
"Africa/Sao_Tome",  
"Africa/Timbuktu",  
"Africa/Tripoli",  
"Africa/Tunis",  
"Africa/Windhoek",  
"America/Adak",  
"America/Anchorage",
```

```
"America/Anguilla",
"America/Antigua",
"America/Araguaina",
"America/Argentina/Buenos_Aires",
"America/Argentina/Catamarca",
"America/Argentina/ComodRivadavia",
"America/Argentina/Cordoba",
"America/Argentina/Jujuy",
"America/Argentina/La_Rioja",
"America/Argentina/Mendoza",
"America/Argentina/Rio_Gallegos",
"America/Argentina/Salta",
"America/Argentina/San_Juan",
"America/Argentina/San_Luis",
"America/Argentina/Tucuman",
"America/Argentina/Ushuaia",
"America/Aruba",
"America/Asuncion",
"America/Atikokan",
"America/Atka",
"America/Bahia",
"America/Bahia_Banderas",
"America/Barbados",
"America/Belem",
"America/Belize",
"America/Blanc-Sablon",
"America/Boa_Vista",
"America/Bogota",
"America/Boise",
"America/Buenos_Aires",
"America/Cambridge_Bay",
"America/Campo_Grande",
"America/Cancun",
"America/Caracas",
"America/Catamarca",
"America/Cayenne",
"America/Cayman",
"America/Chicago",
"America/Chihuahua",
"America/Ciudad_Juarez",
"America/Coral_Harbour",
"America/Cordoba",
"America/Costa_Rica",
"America/Creston",
```

```
"America/Cuiaba",
"America/Curacao",
"America/Danmarkshavn",
"America/Dawson",
"America/Dawson_Creek",
"America/Denver",
"America/Detroit",
"America/Dominica",
"America/Edmonton",
"America/Eirunepe",
"America/El_Salvador",
"America/Ensenada",
"America/Fort_Nelson",
"America/Fort_Wayne",
"America/Fortaleza",
"America/Glace_Bay",
"America/Godthab",
"America/Goose_Bay",
"America/Grand_Turk",
"America/Grenada",
"America/Guadeloupe",
"America/Guatemala",
"America/Guayaquil",
"America/Guyana",
"America/Halifax",
"America/Havana",
"America/Hermosillo",
"America/Indiana/Indianapolis",
"America/Indiana/Knox",
"America/Indiana/Marengo",
"America/Indiana/Petersburg",
"America/Indiana/Tell_City",
"America/Indiana/Vevay",
"America/Indiana/Vincennes",
"America/Indiana/Winamac",
"America/Indianapolis",
"America/Inuvik",
"America/Iqaluit",
"America/Jamaica",
"America/Jujuy",
"America/Juneau",
"America/Kentucky/Louisville",
"America/Kentucky/Monticello",
"America/Knox_IN",
```

```
"America/Kralendijk",
"America/La_Paz",
"America/Lima",
"America/Los_Angeles",
"America/Louisville",
"America/Lower_Princes",
"America/Maceio",
"America/Managua",
"America/Manaus",
"America/Marigot",
"America/Martinique",
"America/Matamoros",
"America/Mazatlan",
"America/Mendoza",
"America/Menominee",
"America/Merida",
"America/Metlakatla",
"America/Mexico_City",
"America/Miquelon",
"America/Moncton",
"America/Monterrey",
"America/Montevideo",
"America/Montreal",
"America/Montserrat",
"America/Nassau",
"America/New_York",
"America/Nipigon",
"America/Nome",
"America/Noronha",
"America/North_Dakota/Beulah",
"America/North_Dakota/Center",
"America/North_Dakota/New_Salem",
"America/Nuuk",
"America/Ojinaga",
"America/Panama",
"America/Pangnirtung",
"America/Paramaribo",
"America/Phoenix",
"America/Port-au-Prince",
"America/Port_of_Spain",
"America/Porto_Acre",
"America/Porto_Velho",
"America/Puerto_Rico",
"America/Punta_Arenas",
```

```
"America/Rainy_River",
"America/Rankin_Inlet",
"America/Recife",
"America/Regina",
"America/Resolute",
"America/Rio_Branco",
"America/Rosario",
"America/Santa_Isabel",
"America/Santarem",
"America/Santiago",
"America/Santo_Domingo",
"America/Sao_Paulo",
"America/Scoresbysund",
"America/Shiprock",
"America/Sitka",
"America/St_Barthelemy",
"America/St_Johns",
"America/St_Kitts",
"America/St_Lucia",
"America/St_Thomas",
"America/St_Vincent",
"America/Swift_Current",
"America/Tegucigalpa",
"America/Thule",
"America/Thunder_Bay",
"America/Tijuana",
"America/Toronto",
"America/Tortola",
"America/Vancouver",
"America/Virgin",
"America/Whitehorse",
"America/Winnipeg",
"America/Yakutat",
"America/Yellowknife",
"Antarctica/Casey",
"Antarctica/Davis",
"Antarctica/DumontDUrville",
"Antarctica/Macquarie",
"Antarctica/Mawson",
"Antarctica/McMurdo",
"Antarctica/Palmer",
"Antarctica/Rothera",
"Antarctica/South_Pole",
"Antarctica/Syowa",
```

```
"Antarctica/Troll",  
"Antarctica/Vostok",  
"Arctic/Longyearbyen",  
"Asia/Aden",  
"Asia/Almaty",  
"Asia/Amman",  
"Asia/Anadyr",  
"Asia/Aqtau",  
"Asia/Aqtobe",  
"Asia/Ashgabat",  
"Asia/Ashkhabad",  
"Asia/Atyrau",  
"Asia/Baghdad",  
"Asia/Bahrain",  
"Asia/Baku",  
"Asia/Bangkok",  
"Asia/Barnaul",  
"Asia/Beirut",  
"Asia/Bishkek",  
"Asia/Brunei",  
"Asia/Calcutta",  
"Asia/Chita",  
"Asia/Choibalsan",  
"Asia/Chongqing",  
"Asia/Chungking",  
"Asia/Colombo",  
"Asia/Dacca",  
"Asia/Damascus",  
"Asia/Dhaka",  
"Asia/Dili",  
"Asia/Dubai",  
"Asia/Dushanbe",  
"Asia/Famagusta",  
"Asia/Gaza",  
"Asia/Harbin",  
"Asia/Hebron",  
"Asia/Ho_Chi_Minh",  
"Asia/Hong_Kong",  
"Asia/Hovd",  
"Asia/Irkutsk",  
"Asia/Istanbul",  
"Asia/Jakarta",  
"Asia/Jayapura",  
"Asia/Jerusalem",
```

```
"Asia/Kabul",  
"Asia/Kamchatka",  
"Asia/Karachi",  
"Asia/Kashgar",  
"Asia/Kathmandu",  
"Asia/Katmandu",  
"Asia/Khandyga",  
"Asia/Kolkata",  
"Asia/Krasnoyarsk",  
"Asia/Kuala_Lumpur",  
"Asia/Kuching",  
"Asia/Kuwait",  
"Asia/Macao",  
"Asia/Macau",  
"Asia/Magadan",  
"Asia/Makassar",  
"Asia/Manila",  
"Asia/Muscat",  
"Asia/Nicosia",  
"Asia/Novokuznetsk",  
"Asia/Novosibirsk",  
"Asia/Omsk",  
"Asia/Oral",  
"Asia/Phnom_Penh",  
"Asia/Pontianak",  
"Asia/Pyongyang",  
"Asia/Qatar",  
"Asia/Qostanay",  
"Asia/Qyzylorda",  
"Asia/Rangoon",  
"Asia/Riyadh",  
"Asia/Saigon",  
"Asia/Sakhalin",  
"Asia/Samarkand",  
"Asia/Seoul",  
"Asia/Shanghai",  
"Asia/Singapore",  
"Asia/Srednekolymsk",  
"Asia/Taipei",  
"Asia/Tashkent",  
"Asia/Tbilisi",  
"Asia/Tehran",  
"Asia/Tel_Aviv",  
"Asia/Thimbu",
```



```
"Asia/Thimphu",  
"Asia/Tokyo",  
"Asia/Tomsk",  
"Asia/Ujung_Pandang",  
"Asia/Ulaanbaatar",  
"Asia/Ulan_Bator",  
"Asia/Urumqi",  
"Asia/Ust-Nera",  
"Asia/Vientiane",  
"Asia/Vladivostok",  
"Asia/Yakutsk",  
"Asia/Yangon",  
"Asia/Yekaterinburg",  
"Asia/Yerevan",  
"Atlantic/Azores",  
"Atlantic/Bermuda",  
"Atlantic/Canary",  
"Atlantic/Cape_Verde",  
"Atlantic/Faeroe",  
"Atlantic/Faroe",  
"Atlantic/Jan_Mayen",  
"Atlantic/Madeira",  
"Atlantic/Reykjavik",  
"Atlantic/South_Georgia",  
"Atlantic/St_Helena",  
"Atlantic/Stanley",  
"Australia/ACT",  
"Australia/Adelaide",  
"Australia/Brisbane",  
"Australia/Broken_Hill",  
"Australia/Canberra",  
"Australia/Currie",  
"Australia/Darwin",  
"Australia/Eucla",  
"Australia/Hobart",  
"Australia/LHI",  
"Australia/Lindeman",  
"Australia/Lord_Howe",  
"Australia/Melbourne",  
"Australia/NSW",  
"Australia/North",  
"Australia/Perth",  
"Australia/Queensland",  
"Australia/South",
```

```
"Australia/Sydney",  
"Australia/Tasmania",  
"Australia/Victoria",  
"Australia/West",  
"Australia/Yancowinna",  
"Brazil/Acre",  
"Brazil/DeNoronha",  
"Brazil/East",  
"Brazil/West",  
"CET",  
"CST6CDT",  
"Canada/Atlantic",  
"Canada/Central",  
"Canada/Eastern",  
"Canada/Mountain",  
"Canada/Newfoundland",  
"Canada/Pacific",  
"Canada/Saskatchewan",  
"Canada/Yukon",  
"Chile/Continental",  
"Chile/EasterIsland",  
"Cuba",  
"EET",  
"EST5EDT",  
"Egypt",  
"Eire",  
"Etc/GMT",  
"Etc/GMT+0",  
"Etc/GMT+1",  
"Etc/GMT+10",  
"Etc/GMT+11",  
"Etc/GMT+12",  
"Etc/GMT+2",  
"Etc/GMT+3",  
"Etc/GMT+4",  
"Etc/GMT+5",  
"Etc/GMT+6",  
"Etc/GMT+7",  
"Etc/GMT+8",  
"Etc/GMT+9",  
"Etc/GMT-0",  
"Etc/GMT-1",  
"Etc/GMT-10",  
"Etc/GMT-11",
```

```
"Etc/GMT-12",  
"Etc/GMT-13",  
"Etc/GMT-14",  
"Etc/GMT-2",  
"Etc/GMT-3",  
"Etc/GMT-4",  
"Etc/GMT-5",  
"Etc/GMT-6",  
"Etc/GMT-7",  
"Etc/GMT-8",  
"Etc/GMT-9",  
"Etc/GMT0",  
"Etc/Greenwich",  
"Etc/UCT",  
"Etc/UTC",  
"Etc/Universal",  
"Etc/Zulu",  
"Europe/Amsterdam",  
"Europe/Andorra",  
"Europe/Astrakhan",  
"Europe/Athens",  
"Europe/Belfast",  
"Europe/Belgrade",  
"Europe/Berlin",  
"Europe/Bratislava",  
"Europe/Brussels",  
"Europe/Bucharest",  
"Europe/Budapest",  
"Europe/Busingen",  
"Europe/Chisinau",  
"Europe/Copenhagen",  
"Europe/Dublin",  
"Europe/Gibraltar",  
"Europe/Guernsey",  
"Europe/Helsinki",  
"Europe/Isle_of_Man",  
"Europe/Istanbul",  
"Europe/Jersey",  
"Europe/Kaliningrad",  
"Europe/Kiev",  
"Europe/Kirov",  
"Europe/Kyiv",  
"Europe/Lisbon",  
"Europe/Ljubljana",
```

```
"Europe/London",  
"Europe/Luxembourg",  
"Europe/Madrid",  
"Europe/Malta",  
"Europe/Mariehamn",  
"Europe/Minsk",  
"Europe/Monaco",  
"Europe/Moscow",  
"Europe/Nicosia",  
"Europe/Oslo",  
"Europe/Paris",  
"Europe/Podgorica",  
"Europe/Prague",  
"Europe/Riga",  
"Europe/Rome",  
"Europe/Samara",  
"Europe/San_Marino",  
"Europe/Sarajevo",  
"Europe/Saratov",  
"Europe/Simferopol",  
"Europe/Skopje",  
"Europe/Sofia",  
"Europe/Stockholm",  
"Europe/Tallinn",  
"Europe/Tirane",  
"Europe/Tiraspol",  
"Europe/Ulyanovsk",  
"Europe/Uzhgorod",  
"Europe/Vaduz",  
"Europe/Vatican",  
"Europe/Vienna",  
"Europe/Vilnius",  
"Europe/Volgograd",  
"Europe/Warsaw",  
"Europe/Zagreb",  
"Europe/Zaporozhye",  
"Europe/Zurich",  
"GB",  
"GB-Eire",  
"GMT",  
"GMT0",  
"Greenwich",  
"Hongkong",  
"Iceland",
```

```
"Indian/Antananarivo",  
"Indian/Chagos",  
"Indian/Christmas",  
"Indian/Cocos",  
"Indian/Comoro",  
"Indian/Kerguelen",  
"Indian/Mahe",  
"Indian/Maldives",  
"Indian/Mauritius",  
"Indian/Mayotte",  
"Indian/Reunion",  
"Iran",  
"Israel",  
"Jamaica",  
"Japan",  
"Kwajalein",  
"Libya",  
"MET",  
"MST7MDT",  
"Mexico/BajaNorte",  
"Mexico/BajaSur",  
"Mexico/General",  
"NZ",  
"NZ-CHAT",  
"Navajo",  
"PRC",  
"PST8PDT",  
"Pacific/Apia",  
"Pacific/Auckland",  
"Pacific/Bougainville",  
"Pacific/Chatham",  
"Pacific/Chuuk",  
"Pacific/Easter",  
"Pacific/Efate",  
"Pacific/Enderbury",  
"Pacific/Fakaofu",  
"Pacific/Fiji",  
"Pacific/Funafuti",  
"Pacific/Galapagos",  
"Pacific/Gambier",  
"Pacific/Guadalcanal",  
"Pacific/Guam",  
"Pacific/Honolulu",  
"Pacific/Johnston",
```

```
"Pacific/Kanton",  
"Pacific/Kiritimati",  
"Pacific/Kosrae",  
"Pacific/Kwajalein",  
"Pacific/Majuro",  
"Pacific/Marquesas",  
"Pacific/Midway",  
"Pacific/Nauru",  
"Pacific/Niue",  
"Pacific/Norfolk",  
"Pacific/Noumea",  
"Pacific/Pago_Pago",  
"Pacific/Palau",  
"Pacific/Pitcairn",  
"Pacific/Pohnpei",  
"Pacific/Ponape",  
"Pacific/Port_Moresby",  
"Pacific/Rarotonga",  
"Pacific/Saipan",  
"Pacific/Samoa",  
"Pacific/Tahiti",  
"Pacific/Tarawa",  
"Pacific/Tongatapu",  
"Pacific/Truk",  
"Pacific/Wake",  
"Pacific/Wallis",  
"Pacific/Yap",  
"Poland",  
"Portugal",  
"ROK",  
"Singapore",  
"SystemV/AST4",  
"SystemV/AST4ADT",  
"SystemV/CST6",  
"SystemV/CST6CDT",  
"SystemV/EST5",  
"SystemV/EST5EDT",  
"SystemV/HST10",  
"SystemV/MST7",  
"SystemV/MST7MDT",  
"SystemV/PST8",  
"SystemV/PST8PDT",  
"SystemV/YST9",  
"SystemV/YST9YDT",
```

```
"Turkey",  
"UCT",  
"US/Alaska",  
"US/Aleutian",  
"US/Arizona",  
"US/Central",  
"US/East-Indiana",  
"US/Eastern",  
"US/Hawaii",  
"US/Indiana-Starke",  
"US/Michigan",  
"US/Mountain",  
"US/Pacific",  
"US/Samoa",  
"UTC",  
"Universal",  
"W-SU",  
"WET",  
"Zulu",  
"EST",  
"HST",  
"MST",  
"ACT",  
"AET",  
"AGT",  
"ART",  
"AST",  
"BET",  
"BST",  
"CAT",  
"CNT",  
"CST",  
"CTT",  
"EAT",  
"ECT",  
"IET",  
"IST",  
"JST",  
"MIT",  
"NET",  
"NST",  
"PLT",  
"PNT",  
"PRT",
```

```
        "PST",
        "SST",
        "VST"
    ]
},
"pageRootPaths": {
    "type": "array",
    "items": {
        "type": "string"
    }
},
"assetRootPaths": {
    "type": "array",
    "items": {
        "type": "string"
    }
},
"crawlAssets": {
    "type": "boolean"
},
"crawlPages": {
    "type": "boolean"
},
"pagePathInclusionPatterns": {
    "type": "array",
    "items": {
        "type": "string"
    }
},
"pagePathExclusionPatterns": {
    "type": "array",
    "items": {
        "type": "string"
    }
},
"pageNameInclusionPatterns": {
    "type": "array",
    "items": {
        "type": "string"
    }
},
"pageNameExclusionPatterns": {
    "type": "array",
    "items": {
```



```
    "type": "string"
  }
},
"assetPathInclusionPatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"assetPathExclusionPatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"assetTypeInclusionPatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"assetTypeExclusionPatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"assetNameInclusionPatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"assetNameExclusionPatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"pageComponents": {
  "type": "array",
  "items": {
    "type": "object"
  }
}
```

```
    },
    "contentFragmentVariations": {
      "type": "array",
      "items": {
        "type": "object"
      }
    },
    "cugExemptedPrincipals": {
      "type": "array",
      "items": {
        "type": "string"
      }
    }
  },
  "required": []
},
"type": {
  "type": "string",
  "pattern": "AEM"
},
"enableIdentityCrawler": {
  "type": "boolean"
},
"syncMode": {
  "type": "string",
  "enum": [
    "FORCED_FULL_CRAWL",
    "FULL_CRAWL",
    "CHANGE_LOG"
  ]
},
"secretArn": {
  "type": "string",
  "minLength": 20,
  "maxLength": 2048
}
},
"version": {
  "type": "string",
  "anyOf": [
    {
      "pattern": "1.0.0"
    }
  ]
}
]
```

```

},
"required": [
  "connectionConfiguration",
  "repositoryConfigurations",
  "syncMode",
  "additionalProperties",
  "secretArn",
  "type"
]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.
repositoryEndpointMetadata	The endpoint information for the data source.
aemUrl	The Adobe Experience Manager host URL. For example, if you use AEM On-Premise, you include the hostname and port: <i>http://hostname:port</i> . Or, if you use AEM as a Cloud Service, you can use the author URL: <i>https://author-xxxxxx-xxxxxxx.adobeaemcloud.com</i> .
authType	The type of authentication you use, whether Basic or OAuth2.
deploymentType	The type of Adobe Experience Manager that you use, either CLOUD or ON-PREMISE .
repositoryConfigurations	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings.
<ul style="list-style-type: none"> page asset 	A list of objects that map the attributes or field names of your Adobe Experience

Configuration	Description
	Manager pages and assets to Amazon Q index field names.
additionalProperties	Additional configuration options for your content in your data source.
isCrawlAcl	Specify true to crawl access control information from documents.
fieldForUserId	Specify field to use for UserId for ACL crawling.
timeZoneId	<p>If you use AEM On-Premise and the time zone of your server is different than the time zone of the Amazon Q AEM connector or index, you can specify the server time zone to align with the AEM connector or index.</p> <p>The default time zone for AEM On-Premise is the time zone of the Amazon Q AEM connector or index. The default time zone for AEM as a Cloud Service is Greenwich Mean Time.</p>
<ul style="list-style-type: none"> • pageRootPaths • assetRootPaths 	A list of root paths for pages and assets. For example, the root path for a page could be <i>/content/sub</i> and the root path for an asset could be <i>/content/sub/asset1</i> .
crawlAssets	Specify true to crawl assets.
crawlPages	Specify true to crawl pages.

Configuration	Description
<ul style="list-style-type: none"> • <code>pagePathInclusionPatterns</code> • <code>pageNameInclusionPatterns</code> • <code>assetPathInclusionPatterns</code> • <code>assetTypeInclusionPatterns</code> • <code>assetNameInclusionPatterns</code> 	<p>A list of regular expression patterns to include certain pages and assets in your Adobe Experience Manager data source. Pages and assets that match the patterns are included in the index. Pages and assets that don't match the patterns are excluded from the index. If a page or asset matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence, and the content isn't included in the index.</p>
<ul style="list-style-type: none"> • <code>pagePathExclusionPatterns</code> • <code>pageNameExclusionPatterns</code> • <code>assetPathExclusionPatterns</code> • <code>assetTypeInclusionPatterns</code> • <code>assetNameInclusionPatterns</code> 	<p>A list of regular expression patterns to exclude certain pages and assets in your Adobe Experience Manager data source. Pages and assets that match the patterns are excluded from the index. Pages and assets that don't match the patterns are included in the index. If a page or asset matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence, and the content isn't included in the index.</p>
<code>pageComponents</code>	<p>A list of names for the specific page components that you want to index.</p>
<code>contentFragmentVariations</code>	<p>A list of names for the specific saved variations of Adobe Experience Manager Content Fragments that you want to index.</p>
<code>type</code>	<p>The type of data source. Specify AEM as your data source type.</p>

Configuration	Description
<code>enableIdentityCrawler</code>	Specify <code>true</code> to use the Amazon Q identity crawler to sync identity/principal information on users and groups with access to specific documents. See Identity crawler for more information.
<code>syncMode</code>	Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose between the following options: <ul style="list-style-type: none">• Use <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index.• Use <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index.• Use <code>CHANGE_LOG</code> to incrementally crawl only new and modified content each time your data source syncs with your index.

Configuration	Description
secretArn	<p>The Amazon Resource Name (ARN) of an AWS Secrets Manager secret that contains the key-value pairs required to connect to your Adobe Experience Manager. The secret must contain a JSON structure with the following keys:</p> <p>If using basic authentication for either AEM On-Premise or Cloud:</p> <pre data-bbox="831 617 1507 974">{ "aemUrl": "<i>Adobe Experience Manager On-Premise host URL</i> ", "username": "<i>user name with admin permissions</i> ", "password": "<i>password with admin permissions</i> " }</pre> <p>If using OAuth 2.0 authentication for AEM On-Premise:</p> <pre data-bbox="831 1129 1507 1444">{ "aemUrl": "<i>Adobe Experience Manager host URL</i>", "clientId": "<i>client ID</i>", "clientSecret": "<i>client secret</i>", "privateKey": "<i>private key</i>" }</pre> <p>If using OAuth 2.0 authentication for AEM as a Cloud Service:</p> <pre data-bbox="831 1600 1507 1820">{ "clientId": "<i>client ID</i>", "clientSecret": "<i>client secret</i>", "privateKey": "<i>private key</i>", "orgId": "<i>organization ID</i> ", }</pre>

Configuration	Description
	<pre data-bbox="829 205 1508 422"> "technicalAccountId": " <i>technical account ID</i>", "imsHost": " <i>Adobe Identity Management System (IMS) host</i> " } </pre>
version	The version of this template that's currently supported.

How Amazon Q Business connector crawls AEM (Server) ACLs

When you connect an AEM (Server) data source to Amazon Q Business, Amazon Q crawls ACL information attached to a document (user and group information) from your AEM (Server) instance. If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

The group and user IDs are mapped as follows:

- `_group_ids` – Group IDs exist in Adobe Experience Manager content where there are set access permissions. They're mapped from the names of the groups in AEM.
- `_user_id` – User IDs exist in Adobe Experience Manager content where there are set access permissions. They're mapped from the user emails as the IDs in AEM.

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business AEM (Server) data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q Business enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a attribute mapping already available, or if you want to map additional document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q Adobe Experience Manager (AEM) connector supports the following entities and the associated reserved and custom attributes.

Supported entities and field mappings

- [Pages](#)
- [Assets](#)

Pages

Amazon Q supports crawling [AEM Pages](#) and offers the following page field mappings.

Adobe Experience Manager (AEM) field name	Index field name	Description	Data type
aem_page_source_uri	_source_uri	Default	String
aem_page_createdBy	_authors	Default	String list
aem_page_template	aem_page_template	Custom	String
aem_entity_type	_category	Default	String
aem_page_createdAt	_created_at	Default	Date
aem_page_lastModified	_last_updated_at	Default	Date
aem_page_lastReplicatedBy	aem_page_publisher	Custom	String
aem_page_lastReplicatedAt	aem_page_publishedAt	Custom	Date

Assets

Amazon Q supports crawling [AEM Assets](#) and offers the following asset field mappings.

Adobe Experience Manager (AEM) field name	Index field name	Description	Data type
aem_page_source_uri	_source_uri	Default	String
aem_page_createdBy	_authors	Default	String list
aem_entity_type	_category	Default	String
aem_page_createdAt	_created_at	Default	Date

Adobe Experience Manager (AEM) field name	Index field name	Description	Data type
aem_page_lastModified	_last_updated_at	Default	Date
aem_page_lastRepliatedBy	aem_page_publisher	Custom	String
aem_page_lastRepliatedAt	aem_page_published At	Custom	Date

IAM role for Amazon Q Business AEM (Server) connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the `BatchPutDocument` and `BatchDeleteDocument` operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
```

```

    "Sid": "AllowsAmazonQToGetSecret",
    "Effect": "Allow",
    "Action": [
      "secretsmanager:GetSecretValue"
    ],
    "Resource": [
      "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToDecryptSecret",
    "Effect": "Allow",
    "Action": [
      "kms:Decrypt"
    ],
    "Resource": [
      "arn:aws:kms:{{region}}:{{account_id}}:key/[key_id]"
    ],
    "Condition": {
      "StringLike": {
        "kms:ViaService": [
          "secretsmanager.*.amazonaws.com"
        ]
      }
    }
  }
},
{
  "Sid": "AllowsAmazonQToIngestDocuments",
  "Effect": "Allow",
  "Action": [
    "qbusiness:BatchPutDocument",
    "qbusiness:BatchDeleteDocument"
  ],
  "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
},
{
  "Sid": "AllowsAmazonQToIngestPrincipalMapping",
  "Effect": "Allow",
  "Action": [
    "qbusiness:PutGroup",
    "qbusiness:CreateUser",
    "qbusiness>DeleteGroup",
    "qbusiness:UpdateUser",

```

```

    "qbusiness:ListGroupsWith",
  ],
  "Resource": [
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}",
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}",
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}/data-source/*"
  ]
},
{
  "Sid": "AllowsAmazonQToCreateAndDeleteNI",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateNetworkInterface",
    "ec2:DeleteNetworkInterface"
  ],
  "Resource": [
    "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[[]subnet_ids[]]",
    "arn:aws:ec2:{{region}}:{{account_id}}:security-group/[[]security_group[]]"
  ]
},
{
  "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateNetworkInterface",
    "ec2:DeleteNetworkInterface"
  ],
  "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
  "Condition": {
    "StringLike": {
      "aws:RequestTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
    },
    "ForAllValues:StringEquals": {
      "aws:TagKeys": [
        "AMAZON_Q"
      ]
    }
  }
},
{
  "Sid": "AllowsAmazonQToCreateTags",
  "Effect": "Allow",

```

```

    "Action": [
      "ec2:CreateTags"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringEquals": {
        "ec2:CreateAction": "CreateNetworkInterface"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterfacePermission"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:ResourceTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
    "Effect": "Allow",
    "Action": [
      "ec2:DescribeNetworkInterfaces",
      "ec2:DescribeAvailabilityZones",
      "ec2:DescribeNetworkInterfaceAttribute",
      "ec2:DescribeVpcs",
      "ec2:DescribeRegions",
      "ec2:DescribeNetworkInterfacePermissions",
      "ec2:DescribeSubnets"
    ],
    "Resource": "*"
  }
]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```
{
```

```

"Version": "2012-10-17",
"Statement": [
  {
    "Sid": "AllowsAmazonQServicePrincipal",
    "Effect": "Allow",
    "Principal": {
      "Service": "qbusiness.amazonaws.com"
    },
    "Action": "sts:AssumeRole",
    "Condition": {
      "StringEquals": {
        "aws:SourceAccount": "{{source_account}}"
      },
      "ArnEquals": {
        "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
      }
    }
  }
]
}

```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Known limitations for the Amazon Q Business AEM (Server) connector

The Amazon Q Business AEM (Server) connector has the following known limitations:

- Deleted site pages can't be tracked when you use **Change log sync** or **Sync only new, modified, or deleted document sync**.

Troubleshooting your Amazon Q Business AEM (Server) connector

The following table provides information about error codes you may see for the Adobe Experience Manager (AEM) connector and suggested troubleshooting actions.

Error code	Error message	Suggested resolution
AEM-5001	Error while getting Administrators group.	Check whether provided username and password are correct or not. Also

Error code	Error message	Suggested resolution
	<p>Below are the possible reasons for this error:</p> <p>Provided AEM host URL might be wrong.</p> <p>Provided username and password are invalid or user is non-admin user.</p>	<p>ensure that the provided user is either admin or belongs to administrators' group.</p>
AEM-5002	Error while generating OAuth2 access token.	Provide valid OAuth2 credentials.
AEM-5103	Null/empty AEM host URL.	AEM host URL should not be null or empty.
AEM-5104	<p>Error while parsing https response. Below are the possible reasons for this error.</p> <ol style="list-style-type: none"> 1. Provided AEM host URL might be wrong, please cross-check the AEM host URL. 2. AEM server is down or not reachable. 	Provide a valid host URL, or try again later.
AEM-5105	Provided authType is incorrect.	Auth type should be Basic or OAuth2.
AEM-5106	Null/empty AEM username.	Username should not be null or empty value.
AEM-5107	Null/empty AEM password.	Password should not be null or empty value.

Error code	Error message	Suggested resolution
AEM-5108	Null/empty client id.	Client Id should not be null or empty value.
AEM-5109	Null/empty client secret.	Client Secret should not be null or empty value.
AEM-5110	Null/empty private key.	Private key should not be null or empty value.
AEM-5111	Null/empty Page Index field name.	Page index field should not be null or empty value
AEM-5112	Null/empty Page data source field name.	Page data source field should not be null or empty value.
AEM-5113	Null/empty asset Index field name.	Asset index field should not be null or empty value.
AEM-5114	Null/empty asset data source field name.	Asset data source field should not be null or empty value.
AEM-5115	Null/empty crawl type.	crawl Type value should be FULL_CRAWL/CHANG_LOG type.
AEM-5116	Invalid AEM host URL format.	Check whether provided AEM URL is in correct format or not e.g. http<s>://<aem-host>:<port>
AEM-5117	Page root paths are incorrect.	Page root paths must be a list of strings.
AEM-5118	Asset root paths are incorrect.	Asset root paths must be a list of strings.
AEM-5119	Page path inclusion or exclusion patterns are incorrect.	Page name inclusion patterns/ Exclusion must be a list of strings.

Error code	Error message	Suggested resolution
AEM-5120	Asset path inclusion or exclusion patterns are incorrect.	Asset name inclusion patterns/ Exclusion must be a list of strings.
AEM-5121	Provided deploymenttype is incorrect.	Deployment type should be either CLOUD or ON_PREMISE.
AEM-5122	Provided orgId is incorrect.	OrgId should not be null or empty value.
AEM-5123	Provided technical Account Id is incorrect.	Technical Account Id should not be null or empty value.
AEM-5124	Provided imsHost is incorrect.	IMS Host should not be null or empty value.
AEM-5125	Null/Empty deployment type.	Deployment type should be either CLOUD or ON_PREMISE.
AEM-5126	Invalid Timezone Id.	Provide a valid timezone id.
AEM-5127	Null/empty asset Index field type.	Asset index field should not be null or empty value.
AEM-5128	Null/empty page Index field type.	Page index field should not be null or empty value.
AEM-5129	DataSourceFieldName doesn't match with IndexFieldType.	Provide a valid asset indexFieldType for the provided asset dataSourceFieldName. Or, provide a valid page indexFieldType for the provided page dataSourceFieldName.
AEM-5130	Protocol used by provided AEM URL is not supported by AEM connector.	Only https protocol is supported by AEM connector. Provide an AEM URL based on https protocol.

Error code	Error message	Suggested resolution
AEM-5131	AEM password is too large.	Password should not be greater than 40 characters.
AEM-5132	AEM client ID is too large.	Client ID should not be greater than 40 characters.
AEM-5133	AEM client secret is too large.	Client secret should not be greater than 40 characters.
AEM-5134	AEM private key is too large.	Private key should not be greater than 2048 characters.
AEM-5135	AEM client ID contains invalid characters.	Client ID should not contain unprintable characters.
AEM-5136	AEM client secret contains invalid characters.	Client secret should not contain unprintable characters.
AEM-5137	AEM private key contains invalid characters.	Private key should not contain unprintable characters.
AEM-5138	AEM IMS host is too large.	IMS host should not be greater than 100 characters.
AEM-5139	AEM technical account ID is too large.	Technical account id should not be greater than 100 characters.
AEM-5140	AEM org ID is too large.	Org id should not be greater than 100 characters.
AEM-5141	Page name inclusion or exclusion patterns are incorrect.	Page name inclusion patterns/ Exclusion must be a list of strings.

Error code	Error message	Suggested resolution
AEM-5142	Asset name inclusion or exclusion patterns are incorrect.	Asset name inclusion patterns/ Exclusion must be a list of strings.
AEM-5143	Asset type inclusion or exclusion patterns are incorrect.	Asset type inclusion patterns/Exclusion must be a list of strings.
AEM-5144	Invalid page root path. Please provide valid page root path.	Page path should start with /content.
AEM-5145	Invalid asset root path. Please provide valid asset root path.	Asset path should start with /content/ dam.
AEM-5146	AEM page root paths list size is too large.	Page root paths list size should not be greater than 1000.
AEM-5147	AEM asset root paths list size is too large.	Asset root paths list size should not be greater than 1000.
AEM-5148	Asset root paths list size should not be greater than 1000.	Asset path exclusion patterns list size should not be greater than 1000.
AEM-5149	AEM asset path inclusion pattern list size is too large.	Asset path inclusion patterns list size should not be greater than 1000.
AEM-5150	AEM asset name inclusion pattern list size is too large.	Asset name inclusion patterns list size should not be greater than 1000.
AEM-5151	AEM asset name exclusion pattern list size is too large.	Asset name exclusion patterns list size should not be greater than 1000.

Error code	Error message	Suggested resolution
AEM-5152	AEM asset type exclusion pattern list size is too large.	Asset type exclusion patterns list size should not be greater than 1000.
AEM-5153	AEM asset type inclusion pattern list size is too large.	Asset type inclusion patterns list size should not be greater than 1000.
AEM-5154	AEM page name inclusion pattern list size is too large.	Page name inclusion patterns list size should not be greater than 1000.
AEM-5155	AEM page name Exclusion pattern list size is too large.	Page name exclusion patterns list size should not be greater than 1000.
AEM-5156	AEM page path Exclusion pattern list size is too large.	Page path exclusion patterns list size should not be greater than 1000.
AEM-5157	AEM page path inclusion pattern list size is too large.	Page path inclusion patterns list size should not be greater than 1000.
AEM-5158	AEM page components list size is too large.	Page components list size should not be greater than 1000.
AEM-5159	AEM content fragment variations list size is too large.	Content fragment variations list size should not be greater than 1000.
AEM-5160	AEM host URL characters length is too large.	AEM host URL characters length should not be greater than 2048 characters.
AEM-5161	Some of the page root paths exceed the character limit.	Page root path characters length should not be greater than 1000.

Error code	Error message	Suggested resolution
AEM-5162	Some of the asset root paths exceed the character limit.	Asset root Path characters length should not be greater than 1000 .
AEM-5163	Some of the asset path exclusion objects exceed the character limit.	Asset path exclusion characters length should not be greater than 1000.
AEM-5164	Some of the asset path inclusion objects exceed the character limit.	Asset path inclusion characters length should not be greater than 1000.
AEM-5165	Some of the asset name inclusion objects exceed the character limit.	Asset name inclusion characters length should not be greater than 1000.
AEM-5166	Some of the asset name exclusion objects exceed the character limit.	Asset name exclusion characters length should not be greater than 1000.
AEM-5167	Some of the asset type exclusion objects exceed the character limit.	Asset type exclusion characters length should not be greater than 1000.
AEM-5168	Some of the asset type inclusion objects exceed the character limit.	Asset type inclusion characters length should not be greater than 1000.
AEM-5169	Some of the page name inclusion objects exceed the character limit.	Page name inclusion characters length should not be greater than 1000.
AEM-5170	Some of the page name exclusion objects exceed the character limit.	Page name exclusion characters length should not be greater than 1000.

Error code	Error message	Suggested resolution
AEM-5171	Some of the page path exclusion objects exceed the character limit.	Page path exclusion characters length should not be greater than 1000.
AEM-5172	Some of the page path inclusion objects exceed the character limit.	Page path inclusion characters length should not be greater than 1000.
AEM-5300	Error in serializing change log token.	Retry sync.
AEM-5301	Error in de-serializing change log token.	Retry sync.
AEM-5401	Error occurred while getting AEM groups.	Retry sync.
AEM-5501	Could not connect to host.	
AEM-5502	AEM URL SSRF check failed.	Make sure AEM host URL is not a multicast/local/link-local/loopback address.
AEM-5503	AEM host not found.	Check whether AEM host is up and reachable.
AEM-5504	Error occurred while executing HTTP request against given AEM URL.	Check whether AEM host is up and reachable.
AEM-5505	AEM malformed URL error.	Provide valid AEM url.
AEM-5506	AEM VPC Configuration check failed.	Site local address is restricted.

Error code	Error message	Suggested resolution
AEM-5507	Error in creating document attribute.	Only String, String List, Date and Long formats are supported for field mappings.
AEM-5200	Error occurred while getting pages from AEM for Full Crawl.	Check whether AEM server is up and responding to API requests.
AEM-5506	AEM VPC Configuration check failed.	Site local address is restricted.
AEM-5201	Error occurred while getting assets from AEM for Full Crawl.	Check whether AEM server is up and responding to API requests.
AEM-5303	Error occurred while getting pages from AEM for Change Log.	Check whether AEM server is up and responding to API requests.
AEM-5304	Error occurred while getting assets from AEM for Change Log.	Check whether AEM server is up and responding to API requests.

Connecting Alfresco (Cloud) to Amazon Q Business

Alfresco is a content management service (CMS) that helps customers store and manage their content. You can connect Alfresco (Cloud) instance to Amazon Q Business—using either the AWS Management Console or the [CreateDataSource](#) API—and create an Amazon Q web experience.

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).
- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [Alfresco \(Cloud\) connector overview](#)
- [Prerequisites for connecting Amazon Q Business to Alfresco \(Cloud\)](#)
- [Connecting Amazon Q Business to Alfresco \(Cloud\) using the console](#)
- [Connecting Amazon Q Business to Alfresco \(Cloud\) using APIs](#)
- [How Amazon Q Business connector crawls Alfresco \(Cloud\) ACLs](#)
- [Amazon Q Business Alfresco \(Cloud\) data source connector field mappings](#)
- [IAM role for Amazon Q Business Alfresco \(Cloud\) connector](#)

Alfresco (Cloud) connector overview

The following table gives an overview of the Amazon Q Business Alfresco (Cloud) connector and its supported features.

Category	Feature	Support
Security	Authentication type	Basic, OAuth 2.0
	Authentication credentials	<p>Basic</p> <ul style="list-style-type: none"> • Alfresco username • Alfresco password <p>OAuth 2.0</p> <ul style="list-style-type: none"> • Client ID • Client secret • Token URL <div style="border: 1px solid #f08080; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p> Important Admin privileges required.</p> </div>

Category	Feature	Support
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Identity crawling	Yes
	VPC	Yes
Crawl features	Custom metadata	Yes
	Entities	Yes. The following entities are supported: <ul style="list-style-type: none"> • Document • Comments
	Field mappings	Yes. Supports both default and custom field mappings. For more information, see Field mappings .
	Filters	Yes. The following filters are supported: <ul style="list-style-type: none"> • Include Aspects • Crawl specific Alfresco site • Include/exclude by file path • Include/exclude by file name • Include/exclude by file type
	Sync mode	Supports full and incremental sync.
	File types	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to Alfresco (Cloud)

Before you begin, make sure that you have completed the following prerequisites.

In Alfresco, make sure you have:

- Copied your Alfresco repository URL and web application URL. If you only want to index a specific Alfresco site, then also copy the site ID.
- Noted your Alfresco authentication credentials, which include a username and password with at least read permissions. If you want to use OAuth 2.0 authentication, you should add the user to the Alfresco administrators group.
- **Optional:** Generated OAuth 2.0 credentials in Alfresco. The credentials include client ID, client secret, and token URL. For more information about how to configure clients for Alfresco On-Premises, see [Alfresco documentation](#). If you use Alfresco Cloud (PaaS), you must contact [Hyland support](#) for Alfresco OAuth 2.0 authentication.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your Alfresco (Cloud) authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Connecting Amazon Q Business to Alfresco (Cloud) using the console

The following procedure outlines how to connect Amazon Q Business to Alfresco (Cloud) using the AWS Management Console.

Connecting Amazon Q to Alfresco (Cloud)

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).

4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **Alfresco (Cloud)** page, enter the following information:

6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. **Source** – Choose **Alfresco Cloud**.

- a. **Alfresco repository URL** – Enter your Alfresco repository URL. For example, if you use Alfresco Cloud (PaaS), the repository URL could be *https://company.alfrescocloud.com*.

- b. **Alfresco user application URL** – Enter your Alfresco user interface URL. You can get the repository URL from your Alfresco administrator. For example, the user interface URL could be *https://example.com*.

8. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

 **Note**

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

9. **Authentication** – Choose **Basic authentication** or **OAuth 2.0 authentication**. Then choose an existing Secrets Manager secret or create a new secret to store your Alfresco credentials. If you choose to create a new secret, an AWS Secrets Manager secret window opens.

If you chose **Basic authentication**, enter a name for the secret, the Alfresco username, and password.

If you chose **OAuth 2.0 authentication**, enter a name for the secret, client ID, client secret, and token URL.

10. **Configure VPC and security group** – *optional* – Choose whether you want to use a VPC. If you do, enter the following information:

- a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
- b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

11. **Identity crawler** – Choose to activate Amazon Q identity crawler to sync identity information. For more information, see [Identity crawler](#).
12. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

13. In **Sync scope**, enter the following information:
 - a. **Content** – Choose whether to crawl content marked with 'Aspects' in Alfresco, content within a specific Alfresco site, or content across all your Alfresco sites.
 - b. **Additional configuration – optional** – Set the following settings:
 - **Include comments** – Choose to include comments in Alfresco Document library and Blog.
 - **Regex patterns** – Regular expression patterns to include or exclude certain files.
14. In **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.
 - **Full sync** – Sync all content regardless of the previous sync status.
 - **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.
15. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
16. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
17. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:

- a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
- b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

 **Note**

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

18. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

19. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

 **Note**

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to Alfresco (Cloud) using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the `configuration` parameter to provide a JSON schema with all other configuration information specific to your data source connector.

For an example of the API request, see [CreateDataSource](#) in the Amazon Q API Reference.

Alfresco JSON schema

The following is the Alfresco JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
      "type": "object",
      "properties": {
        "repositoryEndpointMetadata": {
          "type": "object",
          "properties": {
            "siteId": {
              "type": "string"
            },
            "repoUrl": {
              "type": "string"
            },
            "webAppUrl": {
              "type": "string"
            }
          }
        },
        "repositoryAdditionalProperties": {
          "type": "object",
          "properties": {
            "authType": {
              "type": "string",
              "enum": [
                "OAuth2",
                "Basic"
              ]
            },
            "type": {
              "type": "string",
              "enum": [
                "PAAS",
                "ON_PREM"
              ]
            },
            "crawlType": {
              "type": "string",
              "enum": [
                "ASPECT",
```

```

        "SITE_ID",
        "ALL_SITES"
    ]
  }
}
}
}
},
"required": [
  "repositoryEndpointMetadata"
],
"repositoryConfigurations": {
  "type": "object",
  "properties": {
    "document": {
      "type": "object",
      "properties": {
        "fieldMappings": {
          "type": "array",
          "items": {
            "anyOf": [
              {
                "type": "object",
                "properties": {
                  "indexFieldName": {
                    "type": "string"
                  },
                },
                "indexFieldType": {
                  "type": "string",
                  "enum": [
                    "STRING",
                    "DATE",
                    "STRING_LIST",
                    "LONG"
                  ]
                }
              },
              {
                "dataSourceFieldName": {
                  "type": "string"
                },
                "dateFieldFormat": {
                  "type": "string",
                  "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
                }
              }
            ]
          }
        }
      }
    }
  }
}

```



```

        }
      },
      "required": [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
      ]
    }
  ]
}
},
"required": [
  "fieldMappings"
]
},
"comment": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": {
        "anyOf": [
          {
            "type": "object",
            "properties": {
              "indexFieldName": {
                "type": "string"
              },
            },
            "indexFieldType": {
              "type": "string",
              "enum": [
                "STRING",
                "DATE",
                "STRING_LIST",
                "LONG"
              ]
            },
          },
          {
            "dataSourceFieldName": {
              "type": "string"
            },
          },
          {
            "dateFieldFormat": {
              "type": "string",
              "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
            },
          }
        ]
      }
    }
  }
}

```

```
        }
      },
      "required": [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
      ]
    }
  ]
}
},
"required": [
  "fieldMappings"
]
}
},
"additionalProperties": {
  "type": "object",
  "properties": {
    "isCrawlAcl": {
      "type": "boolean"
    },
    "fieldForUserId": {
      "type": "string"
    },
    "aspectName": {
      "type": "string"
    },
    "aspectProperties": {
      "type": "array"
    },
    "enableFineGrainedControl": {
      "type": "boolean"
    },
    "isCrawlComment": {
      "type": "boolean"
    },
    "inclusionFileNamePatterns": {
      "type": "array"
    },
    "exclusionFileNamePatterns": {
      "type": "array"
    }
  }
}
```

```
    },
    "inclusionFileTypePatterns": {
      "type": "array"
    },
    "exclusionFileTypePatterns": {
      "type": "array"
    },
    "inclusionFilePathPatterns": {
      "type": "array"
    },
    "exclusionFilePathPatterns": {
      "type": "array"
    }
  }
},
"type": {
  "type": "string",
  "pattern": "ALFRESCO"
},
"secretArn": {
  "type": "string",
  "minLength": 20,
  "maxLength": 2048
},
"syncMode": {
  "type": "string",
  "enum": [
    "FORCED_FULL_CRAWL",
    "FULL_CRAWL"
  ]
},
"enableIdentityCrawler": {
  "type": "boolean"
},
"version": {
  "type": "string",
  "anyOf": [
    {
      "pattern": "1.0.0"
    }
  ]
}
},
"required": [
```

```

    "connectionConfiguration",
    "repositoryConfigurations",
    "additionalProperties",
    "type",
    "secretArn"
  ]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.
repositoryEndpointMetadata	The endpoint information for the data source.
siteId	The identifier of the Alfresco site.
repoUrl	The URL of your Alfresco repository. You can get the repository URL from your Alfresco administrator. For example, if you use Alfresco Cloud (PaaS), the repository URL could be <i>https://company.alfrescocloud.com</i> . Or, if you use Alfresco On-Premises, the repository URL could be <i>https://company-alfresco-instance.company-domain.suffix:port</i> .
webAppUrl	The URL of your Alfresco user interface. You can get the Alfresco user interface URL from your Alfresco administrator. For example, the user interface URL could be <i>https://example.com</i> .
repositoryAdditionalProperties	Additional properties for content in your data source.
isCrawlAcl	Specify <code>true</code> to crawl access control information from documents.

Configuration	Description
<code>fieldForUserId</code>	Specify field to use for <code>UserId</code> for ACL crawling.
<code>authType</code>	The type of authentication that you use, whether <code>OAuth2</code> or <code>Basic</code> .
<code>type (deployment)</code>	The type of Alfresco that you use, whether <code>PAAS</code> or <code>ON-PREM</code> .
<code>crawlType</code>	The type of content that you want to crawl, whether <code>ASPECT</code> (content marked with 'Aspects' in Alfresco), <code>SITE_ID</code> (content within a specific Alfresco site), or <code>ALL_SITES</code> (content across all your Alfresco sites).
<code>repositoryConfigurations</code>	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings.
<ul style="list-style-type: none"> document comment 	A list of objects that map the attributes or field names of your Alfresco documents and comments to Amazon Q index field names.
<code>additionalProperties</code>	Additional configuration options for your content in your data source.
<code>aspectProperties</code>	A list of specific 'Aspects' content that you want to index.
<code>enableFineGrainedControl</code>	true to crawl 'Aspects'.
<code>isCrawlComment</code>	true to index comments.

Configuration	Description
<ul style="list-style-type: none">• <code>inclusionFileNamePatterns</code>• <code>inclusionFileTypePatterns</code>• <code>inclusionFilePathPatterns</code>	A list of regular expression patterns to include certain files in your Alfresco data source. Files that match the patterns are included in the index. Files that don't match the patterns are excluded from the index. If a file matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence, and the file isn't included in the index.
<ul style="list-style-type: none">• <code>exclusionFileNamePatterns</code>• <code>exclusionFileTypePatterns</code>• <code>exclusionFilePathPatterns</code>	A list of regular expression patterns to exclude certain files in your Alfresco data source. Files that match the patterns are excluded from the index. Files that don't match the patterns are included in the index. If a file matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence, and the file isn't included in the index.
<code>type</code>	The type of data source. Specify ALFRESCO as your data source type.

Configuration	Description
secretArn	<p>The Amazon Resource Name (ARN) of an AWS Secrets Manager secret that contains the key-value pairs that are required to connect to your Alfresco. The secret must contain a JSON structure with the following keys:</p> <p>If using basic authentication:</p> <pre data-bbox="829 569 1507 768">{ "username": " <i>user name</i>", "password": " <i>password</i>" }</pre> <p>If using OAuth 2.0 authentication:</p> <pre data-bbox="829 877 1507 1115">{ "clientId": " <i>client ID</i>", "clientSecret": " <i>client secret</i>", "tokenUrl": " <i>token URL</i>" }</pre>
syncMode	<p>Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose between the following options:</p> <ul data-bbox="829 1381 1484 1755" style="list-style-type: none">• Use <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index.• Use <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index.

Configuration	Description
<code>enableIdentityCrawler</code>	true to use the Amazon Q identity crawler to sync identity/principal information on users and groups with access to certain documents. See Identity crawler for more information.
<code>version</code>	The version of this template that's currently supported.

How Amazon Q Business connector crawls Alfresco (Cloud) ACLs

When you connect an Alfresco (Cloud) data source to Amazon Q Business, Amazon Q crawls ACL information attached to a document (user and group information) from your Alfresco (Cloud) instance. If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

The group and user IDs are mapped as follows:

- `_group_ids` – Group IDs exist in Alfresco on files where there are set access permissions. They're mapped from the system names of the groups (not display names) in Alfresco.
- `_user_id` – User IDs exist in Alfresco on files where there are set access permissions. They're mapped from the user emails as the IDs in Alfresco.

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business Alfresco (Cloud) data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q Business enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a attribute mapping already available, or if you want to map additional document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q Alfresco connector supports the following entities and the associated reserved and custom attributes.

Supported entities and field mappings

- [Documents](#)
- [Comments](#)

Documents

Alfresco field name	Index field name	Description	Data type
creationTime	_created_at	Default	Date
lastModified	_last_updated_at	Default	Date

Alfresco field name	Index field name	Description	Data type
author	_authors	Default	String list
sourceUri	_source_uri	Default	String
category	_category	Default	String
fileType	_file_type	Default	String
version	_version	Default	String
siteName	al_site_name	Custom	String
size	al_document_size	Custom	Long (numeric)
versionType	al_version_type	Custom	String
title	al_document_title	Custom	String
repositoryId	al_repository_id	Custom	String

Comments

Alfresco field name	Index field name	Description	Data type
creationTime	_created_at	Default	Date
lastModified	_last_updated_at	Default	Date
author	_authors	Default	String list
sourceUri	_source_uri	Default	String
version	_version	Default	String
category	_category	Default	String
fileType	_file_type	Default	String

Alfresco field name	Index field name	Description	Data type
siteName	al_site_name	Custom	String
size	al_document_size	Custom	Long (numeric)
versionType	_al_version_type	Custom	String
title	al_document_title	Custom	String
repositoryId	al_repository_id	Custom	String

IAM role for Amazon Q Business Alfresco (Cloud) connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the `BatchPutDocument` and `BatchDeleteDocument` operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToGetSecret",
      "Effect": "Allow",
```

```

    "Action": [
      "secretsmanager:GetSecretValue"
    ],
    "Resource": [
      "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToDecryptSecret",
    "Effect": "Allow",
    "Action": [
      "kms:Decrypt"
    ],
    "Resource": [
      "arn:aws:kms:{{region}}:{{account_id}}:key/[key_id]"
    ],
    "Condition": {
      "StringLike": {
        "kms:ViaService": [
          "secretsmanager.*.amazonaws.com"
        ]
      }
    }
  }
},
{
  "Sid": "AllowsAmazonQToIngestDocuments",
  "Effect": "Allow",
  "Action": [
    "qbusiness:BatchPutDocument",
    "qbusiness:BatchDeleteDocument"
  ],
  "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
},
{
  "Sid": "AllowsAmazonQToIngestPrincipalMapping",
  "Effect": "Allow",
  "Action": [
    "qbusiness:PutGroup",
    "qbusiness:CreateUser",
    "qbusiness>DeleteGroup",
    "qbusiness:UpdateUser",
    "qbusiness:ListGroup"
  ],
},

```

```

    "Resource": [
      "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}",
      "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}",
      "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}/data-source/*"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateAndDeleteNI",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2:DeleteNetworkInterface"
    ],
    "Resource": [
      "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[subnet_ids]",
      "arn:aws:ec2:{{region}}:{{account_id}}:security-group/[security_group]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2:DeleteNetworkInterface"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:RequestTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
      },
      "ForAllValues:StringEquals": {
        "aws:TagKeys": [
          "AMAZON_Q"
        ]
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateTags",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateTags"
    ]
  }
}

```

```

    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringEquals": {
        "ec2:CreateAction": "CreateNetworkInterface"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterfacePermission"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:ResourceTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
    "Effect": "Allow",
    "Action": [
      "ec2:DescribeNetworkInterfaces",
      "ec2:DescribeAvailabilityZones",
      "ec2:DescribeNetworkInterfaceAttribute",
      "ec2:DescribeVpcs",
      "ec2:DescribeRegions",
      "ec2:DescribeNetworkInterfacePermissions",
      "ec2:DescribeSubnets"
    ],
    "Resource": "*"
  }
]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [

```

```
{
  "Sid": "AllowsAmazonQServicePrincipal",
  "Effect": "Allow",
  "Principal": {
    "Service": "qbusiness.amazonaws.com"
  },
  "Action": "sts:AssumeRole",
  "Condition": {
    "StringEquals": {
      "aws:SourceAccount": "{{source_account}}"
    },
    "ArnEquals": {
      "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
    }
  }
}
```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Connecting Alfresco (Server) to Amazon Q Business

Alfresco is a content management service (CMS) that helps customers store and manage their content. You can connect Alfresco (Server) instance to Amazon Q Business—using either the AWS Management Console or the [CreateDataSource](#) API—and create an Amazon Q web experience.

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).
- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [Alfresco \(Server\) connector overview](#)
- [Prerequisites for connecting Amazon Q Business to Alfresco \(Server\)](#)

- [Connecting Amazon Q Business to Alfresco \(Server\) using the console](#)
- [Connecting Amazon Q Business to Alfresco \(Server\) using APIs](#)
- [How Amazon Q Business connector crawls Alfresco \(Server\) ACLs](#)
- [Amazon Q Business Alfresco \(Server\) data source connector field mappings](#)
- [IAM role for Amazon Q Business Alfresco \(Server\) connector](#)

Alfresco (Server) connector overview

The following table gives an overview of the Amazon Q Business Alfresco (Server) connector and its supported features.

Category	Feature	Support
Security	Authentication type	Basic, OAuth 2.0
	Authentication credentials	<p>Basic</p> <ul style="list-style-type: none"> • Alfresco username • Alfresco password <p>OAuth 2.0</p> <ul style="list-style-type: none"> • Client ID • Client secret • Token URL <div style="border: 1px solid #f08080; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p> Important Admin privileges required.</p> </div>
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Identity crawling	Yes

Category	Feature	Support
	VPC	Yes
Crawl features	Custom metadata	Yes
	Entities	Yes. The following entities are supported: <ul style="list-style-type: none"> • Document • Comments
	Field mappings	Yes. Supports both default and custom field mappings. For more information, see Field mappings .
	Filters	Yes. The following filters are supported: <ul style="list-style-type: none"> • Include Aspects • Crawl specific Alfresco site • Include/exclude by file path • Include/exclude by file name • Include/exclude by file type
	Sync mode	Supports full and incremental sync.
	File types	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to Alfresco (Server)

Before you begin, make sure that you have completed the following prerequisites.

In Alfresco, make sure you have:

- Copied your Alfresco repository URL and web application URL. If you only want to index a specific Alfresco site, then also copy the site ID.
- Noted your Alfresco authentication credentials, which include a username and password with at least read permissions. If you want to use OAuth 2.0 authentication, you should add the user to the Alfresco administrators group.

- **Optional:** Generated OAuth 2.0 credentials in Alfresco. The credentials include client ID, client secret, and token URL. For more information about how to configure clients for Alfresco On-Premises, see [Alfresco documentation](#). If you use Alfresco Cloud (PaaS), you must contact [Hyland support](#) for Alfresco OAuth 2.0 authentication.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your Alfresco (Server) authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Connecting Amazon Q Business to Alfresco (Server) using the console

The following procedure outlines how to connect Amazon Q Business to Alfresco (Server) using the AWS Management Console.

Connecting Amazon Q to Alfresco (Server)

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **Alfresco (Server)** page, enter the following information:

6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. **Source** – Choose **Alfresco server**.

- a. **Alfresco repository URL** – Enter your Alfresco repository URL. For example, if you use Alfresco Cloud (PaaS), the repository URL could be *https://company.alfrescocloud.com*.
- b. **Alfresco user application URL** – Enter your Alfresco user interface URL. You can get the repository URL from your Alfresco administrator. For example, the user interface URL could be *https://example.com*.
- c. **SSL certificate location** – Enter the path to an SSL certificate file stored in an Amazon S3 bucket.

8. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

 **Note**

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

9. **Authentication** – Choose **Basic authentication** or **OAuth 2.0 authentication**. Then choose an existing Secrets Manager secret or create a new secret to store your Alfresco credentials. If you choose to create a new secret, an AWS Secrets Manager secret window opens.

If you chose **Basic authentication**, enter a name for the secret, the Alfresco username, and password.

If you chose **OAuth 2.0 authentication**, enter a name for the secret, client ID, client secret, and token URL.

10. **Configure VPC and security group** – *optional* – Choose whether you want to use a VPC. If you do, enter the following information:

- a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.

- b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

11. **Identity crawler** – Choose to activate Amazon Q identity crawler to sync identity information. For more information, see [Identity crawler](#).
12. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

13. In **Sync scope**, enter the following information:
 - a. **Content** – Choose whether to crawl content marked with 'Aspects' in Alfresco, content within a specific Alfresco site, or content across all your Alfresco sites.
 - b. **Additional configuration – optional** – Set the following settings:
 - **Include comments** – Choose to include comments in Alfresco Document library and Blog.
 - **Regex patterns** – Regular expression patterns to include or exclude certain files.
14. In **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.
 - **Full sync** – Sync all content regardless of the previous sync status.
 - **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.
15. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
16. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
17. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:
 - a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.

- b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

 **Note**

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

18. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

19. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

 **Note**

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to Alfresco (Server) using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the `configuration` parameter to provide a JSON schema with all other configuration information specific to your data source connector.

For an example of the API request, see [CreateDataSource](#) in the Amazon Q API Reference.

Alfresco JSON schema

The following is the Alfresco JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
      "type": "object",
      "properties": {
        "repositoryEndpointMetadata": {
          "type": "object",
          "properties": {
            "siteId": {
              "type": "string"
            },
            "repoUrl": {
              "type": "string"
            },
            "webAppUrl": {
              "type": "string"
            }
          }
        },
        "repositoryAdditionalProperties": {
          "type": "object",
          "properties": {
            "authType": {
              "type": "string",
              "enum": [
                "OAuth2",
                "Basic"
              ]
            },
            "type": {
              "type": "string",
              "enum": [
                "PAAS",
                "ON_PREM"
              ]
            }
          }
        },
        "crawlType": {
          "type": "string",
          "enum": [
            "ASPECT",
            "SITE_ID",
            "ALL_SITES"
          ]
        }
      }
    }
  }
}
```

```

    }
  }
}
},
"required": [
  "repositoryEndpointMetadata"
],
"repositoryConfigurations": {
  "type": "object",
  "properties": {
    "document": {
      "type": "object",
      "properties": {
        "fieldMappings": {
          "type": "array",
          "items": {
            "anyOf": [
              {
                "type": "object",
                "properties": {
                  "indexFieldName": {
                    "type": "string"
                  },
                  "indexFieldType": {
                    "type": "string",
                    "enum": [
                      "STRING",
                      "DATE",
                      "STRING_LIST",
                      "LONG"
                    ]
                  },
                  "dataSourceFieldName": {
                    "type": "string"
                  },
                  "dateFieldFormat": {
                    "type": "string",
                    "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
                  }
                }
              }
            ]
          }
        },
        "required": [

```

```

        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
    ]
  }
]
}
},
"required": [
  "fieldMappings"
]
},
"comment": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": {
        "anyOf": [
          {
            "type": "object",
            "properties": {
              "indexFieldName": {
                "type": "string"
              },
              "indexFieldType": {
                "type": "string",
                "enum": [
                  "STRING",
                  "DATE",
                  "STRING_LIST",
                  "LONG"
                ]
              },
              "dataSourceFieldName": {
                "type": "string"
              },
              "dateFieldFormat": {
                "type": "string",
                "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
              }
            }
          },
          {
            "type": "string"
          }
        ]
      }
    }
  }
},
"required": [

```



```
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
      ]
    }
  ]
}
},
"required": [
  "fieldMappings"
]
}
},
"additionalProperties": {
  "type": "object",
  "properties": {
    "isCrawlAcl": {
      "type": "boolean"
    },
    "fieldForUserId": {
      "type": "string"
    },
    "aspectName": {
      "type": "string"
    },
    "aspectProperties": {
      "type": "array"
    },
    "enableFineGrainedControl": {
      "type": "boolean"
    },
    "isCrawlComment": {
      "type": "boolean"
    },
    "inclusionFileNamePatterns": {
      "type": "array"
    },
    "exclusionFileNamePatterns": {
      "type": "array"
    },
    "inclusionFileTypePatterns": {
      "type": "array"
    }
  }
}
```

```
    },
    "exclusionFileTypePatterns": {
      "type": "array"
    },
    "inclusionFilePathPatterns": {
      "type": "array"
    },
    "exclusionFilePathPatterns": {
      "type": "array"
    }
  }
},
"type": {
  "type": "string",
  "pattern": "ALFRESCO"
},
"secretArn": {
  "type": "string",
  "minLength": 20,
  "maxLength": 2048
},
"syncMode": {
  "type": "string",
  "enum": [
    "FORCED_FULL_CRAWL",
    "FULL_CRAWL"
  ]
},
"enableIdentityCrawler": {
  "type": "boolean"
},
"version": {
  "type": "string",
  "anyOf": [
    {
      "pattern": "1.0.0"
    }
  ]
}
},
"required": [
  "connectionConfiguration",
  "repositoryConfigurations",
  "additionalProperties",
```

```

    "type",
    "secretArn"
  ]
}
```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.
repositoryEndpointMetadata	The endpoint information for the data source.
siteId	The identifier of the Alfresco site.
repoUrl	The URL of your Alfresco repository. You can get the repository URL from your Alfresco administrator. For example, if you use Alfresco Cloud (PaaS), the repository URL could be <i>https://company.alfrescocloud.com</i> . Or, if you use Alfresco On-Premises, the repository URL could be <i>https://company-alfresco-instance.company-domain.suffix:port</i> .
webAppUrl	The URL of your Alfresco user interface. You can get the Alfresco user interface URL from your Alfresco administrator. For example, the user interface URL could be <i>https://example.com</i> .
repositoryAdditionalProperties	Additional properties for content in your data source.
isCrawlAcl	Specify true to crawl access control information from documents.
fieldForUserId	Specify field to use for UserId for ACL crawling.

Configuration	Description
<code>authType</code>	The type of authentication that you use, whether OAuth2 or Basic.
<code>type (deployment)</code>	The type of Alfresco that you use, whether PAAS or ON-PREM.
<code>crawlType</code>	The type of content that you want to crawl, whether ASPECT (content marked with 'Aspects' in Alfresco), SITE_ID (content within a specific Alfresco site), or ALL_SITES (content across all your Alfresco sites).
<code>repositoryConfigurations</code>	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings.
<ul style="list-style-type: none"> • <code>document</code> • <code>comment</code> 	A list of objects that map the attributes or field names of your Alfresco documents and comments to Amazon Q index field names.
<code>additionalProperties</code>	Additional configuration options for your content in your data source.
<code>aspectProperties</code>	A list of specific 'Aspects' content that you want to index.
<code>enableFineGrainedControl</code>	<code>true</code> to crawl 'Aspects'.
<code>isCrawlComment</code>	<code>true</code> to index comments.

Configuration	Description
<ul style="list-style-type: none">• <code>inclusionFileNamePatterns</code>• <code>inclusionFileTypePatterns</code>• <code>inclusionFilePathPatterns</code>	A list of regular expression patterns to include certain files in your Alfresco data source. Files that match the patterns are included in the index. Files that don't match the patterns are excluded from the index. If a file matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence, and the file isn't included in the index.
<ul style="list-style-type: none">• <code>exclusionFileNamePatterns</code>• <code>exclusionFileTypePatterns</code>• <code>exclusionFilePathPatterns</code>	A list of regular expression patterns to exclude certain files in your Alfresco data source. Files that match the patterns are excluded from the index. Files that don't match the patterns are included in the index. If a file matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence, and the file isn't included in the index.
<code>type</code>	The type of data source. Specify ALFRESCO as your data source type.

Configuration	Description
secretArn	<p>The Amazon Resource Name (ARN) of an AWS Secrets Manager secret that contains the key-value pairs that are required to connect to your Alfresco. The secret must contain a JSON structure with the following keys:</p> <p>If using basic authentication:</p> <pre data-bbox="829 569 1507 768">{ "username": " <i>user name</i>", "password": " <i>password</i>" }</pre> <p>If using OAuth 2.0 authentication:</p> <pre data-bbox="829 877 1507 1115">{ "clientId": " <i>client ID</i>", "clientSecret": " <i>client secret</i>", "tokenUrl": " <i>token URL</i>" }</pre>
syncMode	<p>Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose between the following options:</p> <ul data-bbox="829 1381 1485 1753" style="list-style-type: none">• Use <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index.• Use <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index.

Configuration	Description
<code>enableIdentityCrawler</code>	true to use the Amazon Q identity crawler to sync identity/principal information on users and groups with access to certain documents. See Identity crawler for more information.
<code>version</code>	The version of this template that's currently supported.

How Amazon Q Business connector crawls Alfresco (Server) ACLs

When you connect an Alfresco (Server) data source to Amazon Q Business, Amazon Q crawls ACL information attached to a document (user and group information) from your Alfresco (Server) instance. If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

The group and user IDs are mapped as follows:

- `_group_ids` – Group IDs exist in Alfresco on files where there are set access permissions. They're mapped from the system names of the groups (not display names) in Alfresco.
- `_user_id` – User IDs exist in Alfresco on files where there are set access permissions. They're mapped from the user emails as the IDs in Alfresco.

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business Alfresco (Server) data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q Business enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a attribute mapping already available, or if you want to map additional document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q Alfresco connector supports the following entities and the associated reserved and custom attributes.

Supported entities and field mappings

- [Documents](#)
- [Comments](#)

Documents

Alfresco field name	Index field name	Description	Data type
creationTime	_created_at	Default	Date
lastModified	_last_updated_at	Default	Date

Alfresco field name	Index field name	Description	Data type
author	_authors	Default	String list
sourceUri	_source_uri	Default	String
category	_category	Default	String
fileType	_file_type	Default	String
version	_version	Default	String
siteName	al_site_name	Custom	String
size	al_document_size	Custom	Long (numeric)
versionType	al_version_type	Custom	String
title	al_document_title	Custom	String
repositoryId	al_repository_id	Custom	String

Comments

Alfresco field name	Index field name	Description	Data type
creationTime	_created_at	Default	Date
lastModified	_last_updated_at	Default	Date
author	_authors	Default	String list
sourceUri	_source_uri	Default	String
version	_version	Default	String
category	_category	Default	String
fileType	_file_type	Default	String

Alfresco field name	Index field name	Description	Data type
siteName	al_site_name	Custom	String
size	al_document_size	Custom	Long (numeric)
versionType	_al_version_type	Custom	String
title	al_document_title	Custom	String
repositoryId	al_repository_id	Custom	String

IAM role for Amazon Q Business Alfresco (Server) connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the `BatchPutDocument` and `BatchDeleteDocument` operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- Permission to access the SSL certificate stored in your Amazon S3 bucket.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AllowsAmazonQToGetS3Objects",
    "Action": [
```

```

        "s3:GetObject"
    ],
    "Resource": [
        "arn:aws:s3:::{{input_bucket_name}}/*"
    ],
    "Effect": "Allow",
    "Condition": {
        "StringEquals": {
            "aws:ResourceAccount": "{{account_id}}"
        }
    }
},
{
    "Sid": "AllowsAmazonQToGetSecret",
    "Effect": "Allow",
    "Action": [
        "secretsmanager:GetSecretValue"
    ],
    "Resource": [
        "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
    ]
},
{
    "Sid": "AllowsAmazonQToDecryptSecret",
    "Effect": "Allow",
    "Action": [
        "kms:Decrypt"
    ],
    "Resource": [
        "arn:aws:kms:{{region}}:{{account_id}}:key/[key_id]"
    ],
    "Condition": {
        "StringLike": {
            "kms:ViaService": [
                "secretsmanager.*.amazonaws.com"
            ]
        }
    }
},
{
    "Sid": "AllowsAmazonQToIngestDocuments",
    "Effect": "Allow",
    "Action": [
        "qbusiness:BatchPutDocument",

```

```

        "qbusiness:BatchDeleteDocument"
    ],
    "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
  ],
  {
    "Sid": "AllowsAmazonQToIngestPrincipalMapping",
    "Effect": "Allow",
    "Action": [
      "qbusiness:PutGroup",
      "qbusiness:CreateUser",
      "qbusiness>DeleteGroup",
      "qbusiness:UpdateUser",
      "qbusiness:ListGroup"
    ],
    "Resource": [
      "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}",
      "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}",
      "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}/data-source/*"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateAndDeleteNI",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2>DeleteNetworkInterface"
    ],
    "Resource": [
      "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
      "arn:aws:ec2:{{region}}:{{account_id}}:security-group/
[{{security_group}}]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2>DeleteNetworkInterface"
    ],
  },

```

```

    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:RequestTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
      },
      "ForAllValues:StringEquals": {
        "aws:TagKeys": [
          "AMAZON_Q"
        ]
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateTags",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateTags"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringEquals": {
        "ec2:CreateAction": "CreateNetworkInterface"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterfacePermission"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:ResourceTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
    "Effect": "Allow",
    "Action": [

```

```

        "ec2:DescribeNetworkInterfaces",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeNetworkInterfaceAttribute",
        "ec2:DescribeVpcs",
        "ec2:DescribeRegions",
        "ec2:DescribeNetworkInterfacePermissions",
        "ec2:DescribeSubnets"
    ],
    "Resource": "*"
}
]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToAssumeRoleForServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "{{source_account}}"
        },
        "ArnLike": {
          "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
        }
      }
    }
  ]
}

```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Connecting Aurora (MySQL) to Amazon Q Business

Aurora (MySQL) is a relational database management system (RDBMS) built for the cloud. You can connect your Aurora (MySQL) instance to Amazon Q Business—using either the AWS Management Console, CLI, or the [CreateDataSource](#) API—and create an Amazon Q web experience.

The Amazon Q Aurora (MySQL) data source connector supports Aurora MySQL 3 and Aurora Serverless MySQL 8.0.

Important

As a best practice, provide Amazon Q with read-only database credentials. Also, avoid adding tables with sensitive data or personal identifiable information (PII).

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).
- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [Aurora \(MySQL\) connector overview](#)
- [Prerequisites for connecting Amazon Q Business to Aurora \(MySQL\)](#)
- [Connecting Amazon Q Business to Aurora \(MySQL\) using the console](#)
- [Connecting Amazon Q Business to Aurora \(MySQL\) using APIs](#)
- [How Amazon Q Business connector crawls Aurora \(MySQL\) ACLs](#)
- [Amazon Q Business Aurora \(MySQL\) data source connector field mappings](#)
- [IAM role for Amazon Q Business Aurora \(MySQL\) connector](#)
- [Known limitations for the Amazon Q Business Aurora \(MySQL\) connector](#)

Aurora (MySQL) connector overview

The following table gives an overview of the Amazon Q Business Aurora (MySQL) connector and its supported features.

Category	Feature	Support
Security	Authentication type	Basic
	Authentication credentials	<ul style="list-style-type: none"> Username of database user Password of database user
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Driver version	MySQL – 8.0.2.7
	Data source version	Aurora MySQL 3, Aurora Serverless MySQL 8.0
	Identity crawling	No
	VPC	Yes
Crawl features	Custom metadata	Yes
	Entities	<p>Yes. The following entities are supported:</p> <ul style="list-style-type: none"> Document <div style="border: 1px solid #add8e6; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p>Note</p> <p>Each database row is considered an individual searchable Amazon Q document.</p> </div>
	Field mappings	Yes. Supports both default and custom field mappings. For more information, see Field mappings .

Category	Feature	Support
	Sync mode	Supports full and incremental sync.
	File types	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to Aurora (MySQL)

Before you begin, make sure that you have completed the following prerequisites.

In Aurora (MySQL), make sure you have:

- Noted your database user name and password.

Important

As a best practice, provide Amazon Q with read-only database credentials.

- Copied your database host URL, port, and instance. You can find this information on the Amazon RDS console.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your Aurora (MySQL) authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Connecting Amazon Q Business to Aurora (MySQL) using the console

The following procedure outlines how to connect Amazon Q Business to Aurora (MySQL) using the AWS Management Console.

Connecting Amazon Q to Aurora (MySQL)

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **Aurora (MySQL)** page, enter the following information:
6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. In **Source**, enter the following information:
 - a. **Host** – Enter the database host URL, for example: `http://instance URL.region.rds.amazonaws.com`.
 - b. **Port** – Enter the database port, for example, 5432.
 - c. **Instance** – Enter the database instance, for example postgres.
 - d. **SSL certificate location** – Choose to enter the Amazon S3 path to your SSL certificate file.
8. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

Note

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

9. In **Authentication**, enter the following information for your **AWS Secrets Manager secret**.

- a. **Secret name** – A name for your secret.
 - b. For **Database user name**, and **Password** – Enter the authentication credential values you copied from your database.
 - c. Choose **Save**.
10. **Configure VPC and security group – optional** – Choose whether you want to use a VPC. If you do, enter the following information:
- a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
 - b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

11. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

12. In **Sync scope**, enter the following information:
- **SQL query** – Enter SQL query statements like SELECT and JOIN operations. SQL queries must be less than 1000 characters and not contain any semi-colons (;). Amazon Q will crawl all database content that matches your query.
 - **Primary key column** – Provide the primary key for the database table. This identifies a table within your database.
 - **Title column** – Provide the name of the document title column within your database table.
 - **Body column** – Provide the name of the document body column within your database table.
13. In **Additional configuration – optional** – Configure the following settings:
- **Change-detecting columns** – Enter the names of the columns that Amazon Q will use to detect content changes. Amazon Q will re-index content when there is a change in any of these columns.
 - **Users' IDs column** – Enter the name of the column which contains User IDs to be allowed access to content.

- **Groups column** – Enter the name of the column that contains groups to be allowed access to content.
 - **Source URLs column** – Enter the name of the column which contains Source URLs to be indexed.
 - **Time stamps column** – Enter the name of the column which contains time stamps. Amazon Q uses time stamp information to detect changes in your content and sync only changed content.
 - **Time zones column** – Enter the name of the column which contains time zones for the content to be crawled.
 - **Time stamps format** – Enter the name of the column which contains time stamp formats to use to detect content changes and re-sync your content.
14. In **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.
- **Full sync** – Sync all content regardless of the previous sync status.
 - **New or modified content sync** – Sync only new and modified documents.
 - **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.

For more details, see [Sync mode](#).

15. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
16. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
17. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:
- a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
 - b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

Note

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

18. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

19. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

Note

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to Aurora (MySQL) using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the `configuration` parameter to provide a JSON schema with all other configuration information specific to your data source connector.

Aurora (MySQL) JSON schema

The following is the Aurora (MySQL) JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
```

```
"type": "object",
"properties": {
  "repositoryEndpointMetadata": {
    "type": "object",
    "properties": {
      "dbType": {
        "type": "string",
        "enum": [
          "mysql",
          "db2",
          "postgresql",
          "oracle",
          "sqlserver"
        ]
      },
      "dbHost": {
        "type": "string"
      },
      "dbPort": {
        "type": "string"
      },
      "dbInstance": {
        "type": "string"
      }
    },
    "required": [
      "dbType",
      "dbHost",
      "dbPort",
      "dbInstance"
    ]
  },
  "required": [
    "repositoryEndpointMetadata"
  ],
  "repositoryConfigurations": {
    "type": "object",
    "properties": {
      "document": {
        "type": "object",
        "properties": {
          "fieldMappings": {
```

```

    "type": "array",
    "items": [
      {
        "type": "object",
        "properties": {
          "indexFieldName": {
            "type": "string"
          },
          "indexFieldType": {
            "type": "string"
          },
          "dataSourceFieldName": {
            "type": "string"
          }
        },
        "required": [
          "indexFieldName",
          "indexFieldType",
          "dataSourceFieldName"
        ]
      }
    ]
  },
  "required": [
    "fieldMappings"
  ]
},
"required": [
]
},
"additionalProperties": {
  "type": "object",
  "properties": {
    "primaryKey": {
      "type": "string"
    },
    "titleColumn": {
      "type": "string"
    },
    "bodyColumn": {
      "type": "string"
    }
  },

```

```
"sqlQuery": {
  "type": "string",
  "not": {
    "pattern": ";+"
  }
},
"timestampColumn": {
  "type": "string"
},
"timestampFormat": {
  "type": "string"
},
"timezone": {
  "type": "string"
},
"changeDetectingColumns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"allowedUsersColumn": {
  "type": "string"
},
"allowedGroupsColumn": {
  "type": "string"
},
"sourceURIColumn": {
  "type": "string"
},
"serverlessAurora": {
  "type": "string",
  "enum": ["true", "false"]
}
},
"required": ["primaryKey", "titleColumn", "bodyColumn", "sqlQuery"]
},
"type" : {
  "type" : "string",
  "pattern": "JDBC"
},
"syncMode": {
  "type": "string",
  "enum": [
```



```

        "FORCED_FULL_CRAWL",
        "FULL_CRAWL"
    ]
},
"secretArn": {
    "type": "string"
}
},
"version": {
    "type": "string",
    "anyOf": [
        {
            "pattern": "1.0.0"
        }
    ]
}
},
"required": [
    "connectionConfiguration",
    "repositoryConfigurations",
    "syncMode",
    "additionalProperties",
    "secretArn",
    "type"
]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.
repositoryEndpointMetadata	<p>Required configuration information for connecting your data source.</p> <ul style="list-style-type: none"> dbType—The type of Java database you are using, whether <code>mysql</code>, <code>db2</code>, <code>postgresql</code>, <code>oracle</code>, or <code>sqlserver</code>. dbHost—The database host name. dbPort—The database port.

Configuration	Description
	<ul style="list-style-type: none"> • <code>dbInstance</code>—The database instance.
<code>repositoryConfigurations</code>	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings. Specify the type of data source and the secret ARN.
<code>document</code>	A list of objects that map the attributes or field names of your database content to Amazon Q index field names. For more information, see Mapping data source fields .
<code>additionalProperties</code>	Additional configuration options for your content in your data source. Use to include or exclude specific content in your database data source.
<code>primaryKey</code>	Provide the primary key for the database table. This identifies a table within your database.
<code>titleColumn</code>	Provide the name of the document title column within your database table.
<code>bodyColumn</code>	Provide the name of the document title column within your database table.
<code>sqlQuery</code>	Enter SQL query statements like <code>SELECT</code> and <code>JOIN</code> operations. SQL queries must be less than 1000 characters and not contain any semi-colons (;). Amazon Q will crawl all database content that matches your query.

Configuration	Description
timestampColumn	Enter the name of the column which contains time stamps. Amazon Q uses time stamp information to detect changes in your content and sync only changed content.
timestampFormat	Enter the name of the column which contains time stamp formats to use to detect content changes and re-sync your content.
timezone	Enter the name of the column which contains time zones for the content to be crawled.
changeDetectingColumns	Enter the names of the columns that Amazon Q will use to detect content changes. Amazon Q will re-index content when there is a change in any of these columns
allowedUsersColumns	Enter the name of the column which contains User IDs to be allowed access to content.
allowedGroupsColumn	Enter the name of the column which contains User IDs to be allowed access to content.
sourceURIColumn	Enter the name of the column which contains Source URLs to be indexed.
isSslEnabled	<code>true</code> to add a path to an SSL certificate file stored in an Amazon S3 bucket.
type	The type of data source. Specify JDBC as your data source type.

Configuration	Description
syncMode	<p>Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose</p> <ul style="list-style-type: none"> • <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index • <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index • <code>CHANGE_LOG</code> to incrementally crawl only new and modified content each time your data source syncs with your index.
secretArn	<p>The Amazon Resource Name (ARN) of a Secrets Manager secret that contains user name and password required to connect to your database. The secret must contain a JSON structure with the following keys:</p> <pre data-bbox="829 1247 1507 1440"> { "user name": "<i>database user name</i>", "password": "<i>password</i>" } </pre>
version	<p>The version of the template that is currently supported.</p>

How Amazon Q Business connector crawls Aurora (MySQL) ACLs

When you connect a database data source to Amazon Q, Amazon Q crawls user and group information from a column in the source table. You specify this column in the console or using the configuration parameter as part of the `CreateDataSource` operation.

If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

A database data source has the following limitations:

- You can only specify an allow list for a database data source. You can't specify a deny list.
- You can only specify groups. You can't specify individual users for the allow list.
- The database column should be a string containing a semicolon delimited list of groups.

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business Aurora (MySQL) data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a attribute mapping already available, or if you want to map additional document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

⚠ Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q PostgreSQL connector supports the following field mappings:

Supported field mappings

- [Document](#)

Document

JDBC field name	Index field name	Description	Data type
jd_document_id	jd_document_id	Custom	String
jd_document_title	jd_document_title	Custom	String
jd_source_uri	_source_uri	Default	String

IAM role for Amazon Q Business Aurora (MySQL) connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the BatchPutDocument and BatchDeleteDocument operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.

- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- Permission to access the SSL certificate stored in your Amazon S3 bucket.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AllowsAmazonQToGetS3Objects",
    "Action": [
      "s3:GetObject"
    ],
    "Resource": [
      "arn:aws:s3:::{{input_bucket_name}}/*"
    ],
    "Effect": "Allow",
    "Condition": {
      "StringEquals": {
        "aws:ResourceAccount": "{{account_id}}"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToGetSecret",
    "Effect": "Allow",
    "Action": [
      "secretsmanager:GetSecretValue"
    ],
    "Resource": [
      "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToDecryptSecret",
    "Effect": "Allow",
    "Action": [
      "kms:Decrypt"
    ],
    "Resource": [
      "arn:aws:kms:{{region}}:{{account_id}}:key/[key_id]"
    ]
  },

```

```

        "Condition": {
            "StringLike": {
                "kms:ViaService": [
                    "secretsmanager.*.amazonaws.com"
                ]
            }
        },
        {
            "Sid": "AllowsAmazonQToIngestDocuments",
            "Effect": "Allow",
            "Action": [
                "qbusiness:BatchPutDocument",
                "qbusiness:BatchDeleteDocument"
            ],
            "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
        },
        {
            "Sid": "AllowsAmazonQToIngestPrincipalMapping",
            "Effect": "Allow",
            "Action": [
                "qbusiness:PutGroup",
                "qbusiness:CreateUser",
                "qbusiness>DeleteGroup",
                "qbusiness:UpdateUser",
                "qbusiness:ListGroup"
            ],
            "Resource": [
                "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}",
                "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}",
                "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}/data-source/*"
            ]
        },
        {
            "Sid": "AllowsAmazonQToCreateAndDeleteNI",
            "Effect": "Allow",
            "Action": [
                "ec2:CreateNetworkInterface",
                "ec2>DeleteNetworkInterface"
            ],

```



```

        "Resource": [
            "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
            "arn:aws:ec2:{{region}}:{{account_id}}:security-group/
[[security_group]]"
        ]
    },
    {
        "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
        "Effect": "Allow",
        "Action": [
            "ec2:CreateNetworkInterface",
            "ec2:DeleteNetworkInterface"
        ],
        "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
        "Condition": {
            "StringLike": {
                "aws:RequestTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
            },
            "ForAllValues:StringEquals": {
                "aws:TagKeys": [
                    "AMAZON_Q"
                ]
            }
        }
    },
    {
        "Sid": "AllowsAmazonQToCreateTags",
        "Effect": "Allow",
        "Action": [
            "ec2:CreateTags"
        ],
        "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
        "Condition": {
            "StringEquals": {
                "ec2:CreateAction": "CreateNetworkInterface"
            }
        }
    },
    {
        "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
        "Effect": "Allow",
        "Action": [
            "ec2:CreateNetworkInterfacePermission"
        ]
    }
}

```

```

    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:ResourceTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
    "Effect": "Allow",
    "Action": [
      "ec2:DescribeNetworkInterfaces",
      "ec2:DescribeAvailabilityZones",
      "ec2:DescribeNetworkInterfaceAttribute",
      "ec2:DescribeVpcs",
      "ec2:DescribeRegions",
      "ec2:DescribeNetworkInterfacePermissions",
      "ec2:DescribeSubnets"
    ],
    "Resource": "*"
  }
]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToAssumeRoleForServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "{{source_account}}"
        },
        "ArnLike": {

```

```
        "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
    }
}
]
}
```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Known limitations for the Amazon Q Business Aurora (MySQL) connector

- Deleted database rows will not be tracked in when Amazon Q checks for updated content.
- The size of field names and values in a row of your database can't exceed 400KB.
- If you have a large amount of data in your database data source, and do not want Amazon Q to index all your database content after the first sync, you can choose to sync only new, modified, or deleted documents.

Connecting Aurora (PostgreSQL) to Amazon Q Business

Aurora (PostgreSQL) is a relational database management system (RDBMS) built for the cloud. You can connect your Aurora (PostgreSQL) instance to Amazon Q Business—using either the AWS Management Console, CLI, or the [CreateDataSource](#) API—and create an Amazon Q web experience.

The Amazon Q Aurora (PostgreSQL) data source connector supports Aurora PostgreSQL 1.

Important

As a best practice, provide Amazon Q with read-only database credentials. Also, avoid adding tables with sensitive data or personal identifiable information (PII).

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).

- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [Aurora \(PostgreSQL\) connector overview](#)
- [Prerequisites for connecting Amazon Q Business to Aurora \(PostgreSQL\)](#)
- [Connecting Amazon Q Business to Aurora \(PostgreSQL\) using the console](#)
- [Connecting Amazon Q Business to Aurora \(PostgreSQL\) using APIs](#)
- [How Amazon Q Business connector crawls Aurora \(PostgreSQL\) ACLs](#)
- [Amazon Q Business Aurora \(PostgreSQL\) data source connector field mappings](#)
- [IAM role for Amazon Q Business Aurora \(PostgreSQL\) connector](#)
- [Known limitations for the Amazon Q Business Aurora \(PostgreSQL\) connector](#)

Aurora (PostgreSQL) connector overview

The following table gives an overview of the Amazon Q Business Aurora (PostgreSQL) connector and its supported features.

Category	Feature	Support
Security	Authentication type	Basic
	Authentication credentials	<ul style="list-style-type: none"> • Username of database user • Password of database user
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Driver version	Aurora (PostgreSQL) – 42.3.2
	Data source version	Aurora PostgreSQL 1
	Identity crawling	No
	VPC	Yes

Category	Feature	Support
Crawl features	Custom metadata	Yes
	Entities	Yes. The following entities are supported: <ul style="list-style-type: none"> Document <div data-bbox="862 464 1508 730" style="border: 1px solid #add8e6; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p>Note</p> <p>Each database row is considered an individual searchable Amazon Q document.</p> </div>
	Field mappings	Yes. Supports both default and custom field mappings. For more information, see Field mappings .
	Sync mode	Supports full and incremental sync.
	File types	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to Aurora (PostgreSQL)

Before you begin, make sure that you have completed the following prerequisites.

In Aurora (PostgreSQL), make sure you have:

- Noted your database user name and password.

Important

As a best practice, provide Amazon Q with read-only database credentials.

- Copied your database host URL, port, and instance. You can find this information on the Amazon RDS console.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your Aurora (PostgreSQL) authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Connecting Amazon Q Business to Aurora (PostgreSQL) using the console

The following procedure outlines how to connect Amazon Q Business to Aurora (PostgreSQL) using the AWS Management Console.

Connecting Amazon Q to Aurora (PostgreSQL)

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **Aurora (PostgreSQL)** page, enter the following information:
6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. In **Source**, enter the following information:
 - a. **Host** – Enter the database host URL, for example: `http://instance URL.region.rds.amazonaws.com`.

- b. **Port** – Enter the database port, for example, 5432.
 - c. **Instance** – Enter the database instance, for example postgres.
 - d. **Enable SSL certificate location** – Choose to enter the Amazon S3 path to your SSL certificate file.
8. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

 **Note**

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

9. In **Authentication** – Enter the following information for your **AWS Secrets Manager secret**.
 - a. **Secret name** – A name for your secret.
 - b. For **Database user name**, and **Password** – Enter the authentication credential values you copied from your database.
 - c. Choose **Save**.
10. **Configure VPC and security group – optional** – Choose whether you want to use a VPC. If you do, enter the following information:
 - a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
 - b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

11. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

12. In **Sync scope**, enter the following information:

- **SQL query** – Enter SQL query statements like SELECT and JOIN operations. SQL queries must be less than 1000 characters and not contain any semi-colons (;). Amazon Q will crawl all database content that matches your query.
- **Primary key column** – Provide the primary key for the database table. This identifies a table within your database.
- **Title column** – Provide the name of the document title column within your database table.
- **Body column** – Provide the name of the document body column within your database table.

13. In **Additional configuration – optional** – Configure the following settings:

- **Change-detecting columns** – Enter the names of the columns that Amazon Q will use to detect content changes. Amazon Q will re-index content when there is a change in any of these columns.
- **Users' IDs column** – Enter the name of the column which contains User IDs to be allowed access to content.
- **Groups column** – Enter the name of the column that contains groups to be allowed access to content.
- **Source URLs column** – Enter the name of the column which contains Source URLs to be indexed.
- **Time stamps column** – Enter the name of the column which contains time stamps. Amazon Q uses time stamp information to detect changes in your content and sync only changed content.
- **Time zones column** – Enter the name of the column which contains time zones for the content to be crawled.
- **Time stamps format** – Enter the name of the column which contains time stamp formats to use to detect content changes and re-sync your content.

14. In **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.

- **Full sync** – Sync all content regardless of the previous sync status.
- **New or modified content sync** – Sync only new and modified documents.
- **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.

For more details, see [Sync mode](#).

15. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
16. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
17. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:
 - a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
 - b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

 **Note**

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

18. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

19. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

 **Note**

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to Aurora (PostgreSQL) using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the configuration parameter to provide a JSON schema with all other configuration information specific to your data source connector.

Aurora (PostgreSQL) JSON schema

The following is the Aurora (PostgreSQL) JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
      "type": "object",
      "properties": {
        "repositoryEndpointMetadata": {
          "type": "object",
          "properties": {
            "dbType": {
              "type": "string",
              "enum": [
                "mysql",
                "db2",
                "postgresql",
                "oracle",
                "sqlserver"
              ]
            },
            "dbHost": {
              "type": "string"
            },
            "dbPort": {
              "type": "string"
            },
            "dbInstance": {
              "type": "string"
            }
          }
        },
        "required": [
          "dbType",
          "dbHost",
```

```

        "dbPort",
        "dbInstance"
    ]
}
},
"required": [
    "repositoryEndpointMetadata"
]
},
"repositoryConfigurations": {
    "type": "object",
    "properties": {
        "document": {
            "type": "object",
            "properties": {
                "fieldMappings": {
                    "type": "array",
                    "items": [
                        {
                            "type": "object",
                            "properties": {
                                "indexFieldName": {
                                    "type": "string"
                                },
                                "indexFieldType": {
                                    "type": "string"
                                },
                                "dataSourceFieldName": {
                                    "type": "string"
                                }
                            }
                        },
                        {
                            "type": "object",
                            "properties": {
                                "indexFieldName": {
                                    "type": "string"
                                },
                                "indexFieldType": {
                                    "type": "string"
                                },
                                "dataSourceFieldName": {
                                    "type": "string"
                                }
                            }
                        }
                    ]
                }
            }
        }
    },
    "required": [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
    ]
}
}
},
"required": [
    "fieldMappings"
]
}
}

```

```
    },
    "required": [
    ]
  },
  "additionalProperties": {
    "type": "object",
    "properties": {
      "primaryKey": {
        "type": "string"
      },
      "titleColumn": {
        "type": "string"
      },
      "bodyColumn": {
        "type": "string"
      },
      "sqlQuery": {
        "type": "string",
        "not": {
          "pattern": ";+"
        }
      },
      "timestampColumn": {
        "type": "string"
      },
      "timestampFormat": {
        "type": "string"
      },
      "timezone": {
        "type": "string"
      },
      "changeDetectingColumns": {
        "type": "array",
        "items": {
          "type": "string"
        }
      },
      "allowedUsersColumn": {
        "type": "string"
      },
      "allowedGroupsColumn": {
        "type": "string"
      },
      "sourceURIColumn": {
```

```

        "type": "string"
    },
    "serverlessAurora": {
        "type": "string",
        "enum": ["true", "false"]
    }
},
"required": ["primaryKey", "titleColumn", "bodyColumn", "sqlQuery"]
},
"type" : {
    "type" : "string",
    "pattern": "JDBC"
},
"syncMode": {
    "type": "string",
    "enum": [
        "FORCED_FULL_CRAWL",
        "FULL_CRAWL"
    ]
},
"secretArn": {
    "type": "string"
}
},
"version": {
    "type": "string",
    "anyOf": [
        {
            "pattern": "1.0.0"
        }
    ]
}
},
"required": [
    "connectionConfiguration",
    "repositoryConfigurations",
    "syncMode",
    "additionalProperties",
    "secretArn",
    "type"
]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.
repositoryEndpointMetadata	<p>Required configuration information for connecting your data source.</p> <ul style="list-style-type: none"> • dbType—The type of Java database you are using, whether <code>mysql</code>, <code>db2</code>, <code>postgresql</code>, <code>oracle</code>, or <code>sqlserver</code>. • dbHost—The database host name. • dbPort—The database port. • dbInstance—The database instance.
repositoryConfigurations	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings. Specify the type of data source and the secret ARN.
document	A list of objects that map the attributes or field names of your database content to Amazon Q index field names. For more information, see Mapping data source fields .
additionalProperties	Additional configuration options for your content in your data source. Use to include or exclude specific content in your database data source.
primaryKey	Provide the primary key for the database table. This identifies a table within your database.
titleColumn	Provide the name of the document title column within your database table.

Configuration	Description
bodyColumn	Provide the name of the document title column within your database table.
sqlQuery	Enter SQL query statements like SELECT and JOIN operations. SQL queries must be less than 1000 characters and not contain any semi-colons (;). Amazon Q will crawl all database content that matches your query.
timestampColumn	Enter the name of the column which contains time stamps. Amazon Q uses time stamp information to detect changes in your content and sync only changed content.
timestampFormat	Enter the name of the column which contains time stamp formats to use to detect content changes and re-sync your content.
timezone	Enter the name of the column which contains time zones for the content to be crawled.
changeDetectingColumns	Enter the names of the columns that Amazon Q will use to detect content changes. Amazon Q will re-index content when there is a change in any of these columns
allowedUsersColumns	Enter the name of the column which contains User IDs to be allowed access to content.
allowedGroupsColumn	Enter the name of the column which contains User IDs to be allowed access to content.
sourceURIColumn	Enter the name of the column which contains Source URLs to be indexed.
isSslEnabled	true to add a path to an SSL certificate file stored in an Amazon S3 bucket.

Configuration	Description
type	The type of data source. Specify JDBC as your data source type.
syncMode	<p>Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose</p> <ul style="list-style-type: none">• <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index• <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index• <code>CHANGE_LOG</code> to incrementally crawl only new and modified content each time your data source syncs with your index.
secretArn	<p>The Amazon Resource Name (ARN) of a Secrets Manager secret that contains user name and password required to connect to your database. The secret must contain a JSON structure with the following keys:</p> <pre data-bbox="829 1373 1507 1570">{ "user name": "<i>database user name</i>", "password": "<i>password</i>" }</pre>
version	The version of the template that is currently supported.

How Amazon Q Business connector crawls Aurora (PostgreSQL) ACLs

When you connect a database data source to Amazon Q, Amazon Q crawls user and group information from a column in the source table. You specify this column in the console or using the configuration parameter as part of the `CreateDataSource` operation.

If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

A database data source has the following limitations:

- You can only specify an allow list for a database data source. You can't specify a deny list.
- You can only specify groups. You can't specify individual users for the allow list.
- The database column should be a string containing a semicolon delimited list of groups.

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business Aurora (PostgreSQL) data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a attribute mapping already available, or if you want to map additional

document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q PostgreSQL connector supports the following field mappings:

Supported field mappings

- [Document](#)

Document

JDBC field name	Index field name	Description	Data type
jd_document_id	jd_document_id	Custom	String
jd_document_title	jd_document_title	Custom	String
jd_source_uri	_source_uri	Default	String

IAM role for Amazon Q Business Aurora (PostgreSQL) connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the `BatchPutDocument` and `BatchDeleteDocument` operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- Permission to access the SSL certificate stored in your Amazon S3 bucket.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AllowsAmazonQToGetS3Objects",
    "Action": [
      "s3:GetObject"
    ],
    "Resource": [
      "arn:aws:s3:::{{input_bucket_name}}/*"
    ],
    "Effect": "Allow",
    "Condition": {
      "StringEquals": {
        "aws:ResourceAccount": "{{account_id}}"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToGetSecret",
    "Effect": "Allow",
    "Action": [
      "secretsmanager:GetSecretValue"
    ],
    "Resource": [
      "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToDecryptSecret",
    "Effect": "Allow",
    "Action": [
```

```

        "kms:Decrypt"
    ],
    "Resource": [
        "arn:aws:kms:{{region}}:{{account_id}}:key/[key_id]"
    ],
    "Condition": {
        "StringLike": {
            "kms:ViaService": [
                "secretsmanager.*.amazonaws.com"
            ]
        }
    }
},
{
    "Sid": "AllowsAmazonQToIngestDocuments",
    "Effect": "Allow",
    "Action": [
        "qbusiness:BatchPutDocument",
        "qbusiness:BatchDeleteDocument"
    ],
    "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
},
{
    "Sid": "AllowsAmazonQToIngestPrincipalMapping",
    "Effect": "Allow",
    "Action": [
        "qbusiness:PutGroup",
        "qbusiness:CreateUser",
        "qbusiness>DeleteGroup",
        "qbusiness:UpdateUser",
        "qbusiness:ListGroups"
    ],
    "Resource": [
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}",
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}",
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}/data-source/*"
    ]
},
{
    "Sid": "AllowsAmazonQToCreateAndDeleteNI",

```

```

    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2>DeleteNetworkInterface"
    ],
    "Resource": [
      "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
      "arn:aws:ec2:{{region}}:{{account_id}}:security-group/
[[security_group]]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2>DeleteNetworkInterface"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:RequestTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
      },
      "ForAllValues:StringEquals": {
        "aws:TagKeys": [
          "AMAZON_Q"
        ]
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateTags",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateTags"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringEquals": {
        "ec2:CreateAction": "CreateNetworkInterface"
      }
    }
  }
},

```

```

    {
      "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
      "Effect": "Allow",
      "Action": [
        "ec2:CreateNetworkInterfacePermission"
      ],
      "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
      "Condition": {
        "StringLike": {
          "aws:ResourceTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
        }
      }
    },
    {
      "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
      "Effect": "Allow",
      "Action": [
        "ec2:DescribeNetworkInterfaces",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeNetworkInterfaceAttribute",
        "ec2:DescribeVpcs",
        "ec2:DescribeRegions",
        "ec2:DescribeNetworkInterfacePermissions",
        "ec2:DescribeSubnets"
      ],
      "Resource": "*"
    }
  ]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToAssumeRoleForServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
    }
  ]
}

```

```
    "Condition": {
      "StringEquals": {
        "aws:SourceAccount": "{{source_account}}"
      },
      "ArnLike": {
        "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
      }
    }
  }
]
```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Known limitations for the Amazon Q Business Aurora (PostgreSQL) connector

- Deleted database rows will not be tracked in when Amazon Q checks for updated content.
- The size of field names and values in a row of your database can't exceed 400KB.
- If you have a large amount of data in your database data source, and do not want Amazon Q to index all your database content after the first sync, you can choose to sync only new, modified, or deleted documents.

Connecting Amazon FSx (Windows) to Amazon Q Business

Amazon FSx (Windows) is a fully managed, cloud based file server system that offers shared storage capabilities. You can connect your Amazon FSx (Windows) instance to Amazon Q Business—using either the AWS Management Console, CLI, or the [CreateDataSource](#) API—and create an Amazon Q web experience.

The Amazon Q Amazon FSx (Windows) data source connector supports only Amazon FSx for Windows.

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).

- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [Amazon FSx \(Windows\) connector overview](#)
- [Prerequisites for connecting Amazon Q Business to Amazon FSx \(Windows\)](#)
- [Connecting Amazon Q Business to Amazon FSx \(Windows\) using the console](#)
- [Connecting Amazon Q Business to Amazon FSx \(Windows\) using APIs](#)
- [How Amazon Q Business connector crawls Amazon FSx \(Windows\) ACLs](#)
- [Amazon Q Business Amazon FSx \(Windows\) data source connector field mappings](#)
- [IAM role for Amazon Q Business Amazon FSx \(Windows\) connector](#)

Amazon FSx (Windows) connector overview

The following table gives an overview of the Amazon Q Business Amazon FSx (Windows) connector and its supported features.

Category	Feature	Support
Security	Authentication type	Basic
	Authentication credentials	<ul style="list-style-type: none"> • Amazon FSx (Windows) username • Amazon FSx (Windows) password
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Identity crawling	No
	VPC	Yes
Crawl features	Custom metadata	Yes
	Field mappings	Yes. Supports both default and custom field mappings. For more information, see Field mappings .

Category	Feature	Support
	Filters	Yes. The following filters are supported: <ul style="list-style-type: none"> • Include/exclude by file name • Include/exclude by file type • Include/exclude by file path
	Sync mode	Supports full and incremental sync.
	File types	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to Amazon FSx (Windows)

Before you begin, make sure that you have completed the following prerequisites.

In Amazon FSx (Windows), make sure you have:

- An Amazon FSx (Windows) account with read and mounting permissions.
- Noted your Amazon FSx authentication credentials for an Active Directory user account. This includes your Active Directory user name and your Domain Name System (DNS) domain name. For example, *user@corp.example.com*.
- Copied your Amazon FSx file system ID.
- Used an Amazon VPC (AWS VPC) where your Amazon FSx resides.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your Amazon FSx (Windows) authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Connecting Amazon Q Business to Amazon FSx (Windows) using the console

The following procedure outlines how to connect Amazon Q Business to Amazon FSx (Windows) using the AWS Management Console.

Connecting Amazon Q to Amazon FSx (Windows)

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **Amazon FSx (Windows)** page, enter the following information:
6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. In **Source**, for **Amazon FSx file system ID**—Select your file system ID or create a new directory.

Only already created file system IDs are displayed and available to connect.

8. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

Note

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

9. **Authentication** – Enter the following information for your **AWS Secrets Manager secret**.
 - a. **Secret name**—A name for your secret.

- b. For **User name**—Enter the user name for Amazon FSx Active Directory account.
 - c. For **Password**—Enter the password for the Amazon FSx Active Directory account.
 - d. Choose **Save**.
10. **Configure VPC and security group – optional** – Choose whether you want to use a VPC. If you do, enter the following information:
- a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
 - b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

11. **Identity crawler** – Choose to activate Amazon Q identity crawler to sync identity information. For more information, see [Identity crawler](#).
12. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

13. In **Sync scope**, enter the following information:
- **Regex patterns**—Add regular expression patterns to include or exclude certain content. You can add up to 100 patterns.
14. In **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.
- **Full sync** – Sync all content regardless of the previous sync status.
 - **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.

For more details, see [Sync mode](#).

15. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
16. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.

17. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:
- Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
 - Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

 **Note**

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

18. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

19. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

 **Note**

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to Amazon FSx (Windows) using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the `configuration` parameter to provide a JSON schema with all other configuration information specific to your data source connector.

For an example of the API request, see [CreateDataSource](#) in the Amazon Q API Reference.

Amazon FSx JSON schema

The following is the Amazon FSx JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
      "type": "object",
      "properties": {
        "repositoryEndpointMetadata": {
          "type": "object",
          "properties": {
            "fileSystemId": {
              "type": "string",
              "pattern": "fs-.*"
            },
            "fileSystemType": {
              "type": "string",
              "pattern": "WINDOWS"
            }
          },
          "required": ["fileSystemId", "fileSystemType"]
        }
      }
    },
    "repositoryConfigurations": {
      "type": "object",
      "properties": {
        "All": {
          "type": "object",
          "properties": {
            "fieldMappings": {
              "type": "array",
              "items": [
                {
                  "type": "object",
                  "properties": {
                    "indexFieldName": {
                      "type": "string"
                    }
                  }
                }
              ]
            }
          }
        }
      }
    }
  }
}
```

```

        "indexFieldType": {
            "type": "string",
            "enum": ["STRING", "STRING_LIST", "DATE"]
        },
        "dataSourceFieldName": {
            "type": "string"
        },
        "dateFieldFormat": {
            "type": "string",
            "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
        }
    },
    "required": [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
    ]
}
]
}
},
"required": ["fieldMappings"]
}
},
"required": ["All"]
},
"additionalProperties": {
    "type": "object",
    "properties": {
        "isCrawlAcl": {
            "type": "boolean"
        },
        "exclusionPatterns": {
            "type": "array",
            "items": {
                "type": "string"
            }
        },
        "inclusionPatterns": {
            "type": "array",
            "items": {
                "type": "string"
            }
        }
    }
}
}

```

```

    },
    "required": []
  },
  "enableIdentityCrawler": {
    "type": "boolean"
  },
  "syncMode": {
    "type": "string",
    "enum": [
      "FORCED_FULL_CRAWL",
      "FULL_CRAWL"
    ]
  },
  "type" : {
    "type" : "string",
    "pattern": "FSX"
  }
},
"version": {
  "type": "string",
  "anyOf": [
    {
      "pattern": "1.0.0"
    }
  ]
},
"required": [
  "connectionConfiguration",
  "repositoryConfigurations",
  "syncMode",
  "enableIdentityCrawler",
  "additionalProperties",
  "type"
]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.

Configuration	Description
repositoryEndpointMetadata	The endpoint information for the data source.
fileSystemId	The identifier of the Amazon FSx (Windows) file system. You can find your file system ID on the File Systems dashboard in the Amazon FSx (Windows) console.
fileSystemType	The type of Amazon FSx you use: ONTAP.
repositoryConfigurations	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings.
<ul style="list-style-type: none"> All 	A list of objects that map the attributes or field names of your Amazon FSx (Windows) pages and assets to Amazon Q index field names.
additionalProperties	Additional configuration options for your content in your data source.
<ul style="list-style-type: none"> inclusionPatterns 	A list of regular expression patterns to include specific content from you Amazon FSx (Windows) data source. Content that match the patterns are included in the index. Content that doesn't match the patterns are excluded from the index. If content matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence, and the content isn't included in the index.

Configuration	Description
<ul style="list-style-type: none"> <code>exclusionPatterns</code> 	<p>A list of regular expression patterns to exclude specific content from your Amazon FSx (Windows) data source. Content that match the patterns are excluded from the index. Content that doesn't match the patterns are included in the index. If content matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence, and the content isn't included in the index.</p>
<code>enableIdentityCrawler</code>	<p><code>true</code> to activate identity crawler. Identity crawler is activated by default. Crawling identity information on users and groups with access to specific documents is useful for user context filtering. Search results are filtered based on the user or their group access to documents. See Identity crawler for more information.</p>
<code>syncMode</code>	<p>Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose between the following options:</p> <ul style="list-style-type: none"> Use <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index Use <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index
<code>type</code>	<p>The type of data source. Specify FSX as your data source type.</p>

Configuration	Description
version	The version of this template that's currently supported.

How Amazon Q Business connector crawls Amazon FSx (Windows) ACLs

When you connect an Amazon FSx (Windows) data source to Amazon Q Business, Amazon Q Business crawls ACL information attached to a document (user and group information) from the directory service of the Amazon FSx instance. If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

The group and user IDs are mapped as follows:

- `_group_ids`—Group IDs exist in Amazon FSx on files where there are set access permissions. They are mapped from the system group names in the directory service of Amazon FSx.
- `_user_id`—User IDs exist in Amazon FSx on files where there are set access permissions. They are mapped from the system user names in the directory service of Amazon FSx.

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business Amazon FSx (Windows) data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q Business enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a attribute mapping already available, or if you want to map additional document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

⚠ Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q Amazon FSx (Windows) connector supports the following entities and the associated reserved and custom attributes.

Amazon FSx (Windows) field name	Index field name	Description	Data type
creationTime	_created_at	Default	Date
lastModified	_last_updated_at	Default	Date
fileType	_file_type	Default	String
path	_source_uri	Default	String
author	_authors	Default	String list
size	fsx_size	Custom	String
lastAccessTime	_last_accessed_at	Custom	Date

IAM role for Amazon Q Business Amazon FSx (Windows) connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the `BatchPutDocument` and `BatchDeleteDocument` operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- Permission to access the SSL certificate stored in your Amazon S3 bucket.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AllowsAmazonQToGetS3Objects",
    "Action": [
      "s3:GetObject"
    ],
    "Resource": [
      "arn:aws:s3:::{{input_bucket_name}}/*"
    ],
    "Effect": "Allow",
    "Condition": {
      "StringEquals": {
        "aws:ResourceAccount": "{{account_id}}"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToGetSecret",
```

```

    "Effect": "Allow",
    "Action": [
      "secretsmanager:GetSecretValue"
    ],
    "Resource": [
      "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[{{secret_id}}]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToDecryptSecret",
    "Effect": "Allow",
    "Action": [
      "kms:Decrypt"
    ],
    "Resource": [
      "arn:aws:kms:{{region}}:{{account_id}}:key/[{{key_id}}]"
    ],
    "Condition": {
      "StringLike": {
        "kms:ViaService": [
          "secretsmanager.*.amazonaws.com"
        ]
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToIngestDocuments",
    "Effect": "Allow",
    "Action": [
      "qbusiness:BatchPutDocument",
      "qbusiness:BatchDeleteDocument"
    ],
    "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
    {{application_id}}/index/{{index_id}}"
  },
  {
    "Sid": "AllowsAmazonQToIngestPrincipalMapping",
    "Effect": "Allow",
    "Action": [
      "qbusiness:PutGroup",
      "qbusiness:CreateUser",
      "qbusiness>DeleteGroup",
      "qbusiness:UpdateUser",
      "qbusiness:ListGroup"
    ]
  }
}

```

```

    ],
    "Resource": [
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}",
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}",
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}/data-source/*"
    ]
},
{
    "Sid": "AllowsAmazonQToCreateAndDeleteNI",
    "Effect": "Allow",
    "Action": [
        "ec2:CreateNetworkInterface",
        "ec2:DeleteNetworkInterface"
    ],
    "Resource": [
        "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
        "arn:aws:ec2:{{region}}:{{account_id}}:security-group/
[{{security_group}}]"
    ]
},
{
    "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
    "Effect": "Allow",
    "Action": [
        "ec2:CreateNetworkInterface",
        "ec2:DeleteNetworkInterface"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
        "StringLike": {
            "aws:RequestTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
        },
        "ForAllValues:StringEquals": {
            "aws:TagKeys": [
                "AMAZON_Q"
            ]
        }
    }
},
{

```

```

        "Sid": "AllowsAmazonQToCreateTags",
        "Effect": "Allow",
        "Action": [
            "ec2:CreateTags"
        ],
        "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
        "Condition": {
            "StringEquals": {
                "ec2:CreateAction": "CreateNetworkInterface"
            }
        }
    },
    {
        "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
        "Effect": "Allow",
        "Action": [
            "ec2:CreateNetworkInterfacePermission"
        ],
        "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
        "Condition": {
            "StringLike": {
                "aws:ResourceTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
            }
        }
    },
    {
        "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
        "Effect": "Allow",
        "Action": [
            "ec2:DescribeNetworkInterfaces",
            "ec2:DescribeAvailabilityZones",
            "ec2:DescribeNetworkInterfaceAttribute",
            "ec2:DescribeVpcs",
            "ec2:DescribeRegions",
            "ec2:DescribeNetworkInterfacePermissions",
            "ec2:DescribeSubnets"
        ],
        "Resource": "*"
    }
]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToAssumeRoleForServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "{{source_account}}"
        },
        "ArnLike": {
          "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
        }
      }
    }
  ]
}
```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Connecting Amazon RDS (Microsoft SQL Server) to Amazon Q Business

Amazon RDS (Microsoft SQL Server) is a relational database management system (RDBMS) built for the cloud. You can connect your Amazon RDS (Microsoft SQL Server) instance to Amazon Q Business – using either the AWS Management Console, CLI, or the [CreateDataSource](#) API – and create an Amazon Q web experience.

The Amazon Q Microsoft SQL Server data source connector supports MS SQL Server 2019.

Important

As a best practice, provide Amazon Q with read-only database credentials. Also, avoid adding tables with sensitive data or personal identifiable information (PII).

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).
- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [Amazon RDS \(Microsoft SQL Server\) connector overview](#)
- [Prerequisites for connecting Amazon Q Business to Amazon RDS \(Microsoft SQL Server\)](#)
- [Connecting Amazon Q Business to Amazon RDS \(Microsoft SQL Server\) using the console](#)
- [Connecting Amazon Q Business to Amazon RDS \(Microsoft SQL Server\) using APIs](#)
- [How Amazon Q Business connector crawls Amazon RDS \(Microsoft SQL Server\) ACLs](#)
- [Amazon Q Business Amazon RDS \(Microsoft SQL Server\) data source connector field mappings](#)
- [IAM role for Amazon Q Business Amazon RDS \(Microsoft SQL Server\) connector](#)
- [Known limitations for the Amazon Q Business Amazon RDS \(Microsoft SQL Server\) connector](#)

Amazon RDS (Microsoft SQL Server) connector overview

The following table gives an overview of the Amazon Q Business Amazon RDS (Microsoft SQL Server) connector and its supported features.

Category	Feature	Support
Security	Authentication type	Basic
	Authentication credentials	<ul style="list-style-type: none"> • Username of database user • Password of database user
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Driver version	Microsoft SQL Server – 10.2.0.jre11

Category	Feature	Support
	Data source version	Microsoft SQL Server 2019
	Identity crawling	No
	VPC	Yes
Crawl features	Custom metadata	Yes
	Entities	Yes. The following entities are supported: <ul style="list-style-type: none"> Document <div data-bbox="862 705 1508 974" style="border: 1px solid #add8e6; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p> Note Each database row is considered an individual searchable Amazon Q document.</p> </div>
	Field mappings	Yes. Supports both default and custom field mappings. For more information, see Field mappings .
	Sync mode	Supports full and incremental sync.
	File types	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to Amazon RDS (Microsoft SQL Server)

Before you begin, make sure that you have completed the following prerequisites.

In Amazon RDS (Microsoft SQL Server), make sure you have:

- Noted your database user name and password.

⚠ Important

As a best practice, provide Amazon Q with read-only database credentials.

- Copied your database host URL, port, and instance.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your Amazon RDS (Microsoft SQL Server) authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

📘 Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Connecting Amazon Q Business to Amazon RDS (Microsoft SQL Server) using the console

On the **Amazon RDS (Microsoft SQL Server)** page, enter the following information:

1. **Name** – Name your data source for easy tracking.
Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.
2. In **Source**, enter the following information:
 - a. **Host** – Enter the database host name.
 - b. **Port** – Enter the database port.
 - c. **Instance** – Enter the database instance.
 - d. **SSL certificate location** – Choose to enter the Amazon S3 path to your SSL certificate file.
3. In **Authentication** – Enter the following information for your **AWS Secrets Manager secret**.

- a. **Secret name** – A name for your secret.
 - b. For **Database user name**, and **Password** – Enter the authentication credential values you copied from your database.
 - c. Choose **Save**.
4. **Configure VPC and security group – optional** – Choose whether you want to use a VPC. If you do, enter the following information:
- a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
 - b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

5. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

6. In **Sync scope**, enter the following information:
- **SQL query** – Enter SQL query statements like SELECT and JOIN operations. SQL queries must be less than 1000 characters and not contain any semi-colons (;). Amazon Q will crawl all database content that matches your query.

If a table name has special characters, put it in square brackets "[]" in the SQL query. For example: `select * from [my-database-table]`.

- **Primary key column** – Provide the primary key for the database table. This identifies a table within your database.
 - **Title column** – Provide the name of the document title column within your database table.
 - **Body column** – Provide the name of the document body column within your database table.
7. In **Additional configuration – optional** – Configure the following settings:
- **Change-detecting columns** – Enter the names of the columns that Amazon Q will use to detect content changes. Amazon Q will re-index content when there is a change in any of these columns.

- **Users' IDs column** – Enter the name of the column which contains User IDs to be allowed access to content.
 - **Groups column** – Enter the name of the column that contains groups to be allowed access to content.
 - **Source URLs column** – Enter the name of the column which contains Source URLs to be indexed.
 - **Time stamps column** – Enter the name of the column which contains time stamps. Amazon Q uses time stamp information to detect changes in your content and sync only changed content.
 - **Time zones column** – Enter the name of the column which contains time zones for the content to be crawled.
 - **Time stamps format** – Enter the name of the column which contains time stamp formats to use to detect content changes and re-sync your content.
8. In **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.
- **Full sync** – Sync all content regardless of the previous sync status.
 - **New or modified content sync** – Sync only new and modified documents.
 - **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.

For more details, see [Sync mode](#).

9. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
10. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
11. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:
- Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
 - Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

Note

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

12. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

13. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

Note

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to Amazon RDS (Microsoft SQL Server) using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q Business application.

Then, you use the configuration parameter to provide a JSON schema with all other configuration information specific to your data source connector.

Amazon RDS (Microsoft SQL Server) JSON schema

The following is the Amazon RDS (Microsoft SQL Server) JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
```

```
"connectionConfiguration": {
  "type": "object",
  "properties": {
    "repositoryEndpointMetadata": {
      "type": "object",
      "properties": {
        "dbType": {
          "type": "string",
          "enum": [
            "mysql",
            "db2",
            "postgresql",
            "oracle",
            "sqlserver"
          ]
        },
        "dbHost": {
          "type": "string"
        },
        "dbPort": {
          "type": "string"
        },
        "dbInstance": {
          "type": "string"
        }
      },
      "required": [
        "dbType",
        "dbHost",
        "dbPort",
        "dbInstance"
      ]
    },
    "required": [
      "repositoryEndpointMetadata"
    ]
  },
  "repositoryConfigurations": {
    "type": "object",
    "properties": {
      "document": {
        "type": "object",
        "properties": {
```

```

    "fieldMappings": {
      "type": "array",
      "items": [
        {
          "type": "object",
          "properties": {
            "indexFieldName": {
              "type": "string"
            },
            "indexFieldType": {
              "type": "string"
            },
            "dataSourceFieldName": {
              "type": "string"
            }
          },
          "required": [
            "indexFieldName",
            "indexFieldType",
            "dataSourceFieldName"
          ]
        }
      ]
    },
    "required": [
      "fieldMappings"
    ]
  },
  "required": [
  ],
  "additionalProperties": {
    "type": "object",
    "properties": {
      "primaryKey": {
        "type": "string"
      },
      "titleColumn": {
        "type": "string"
      },
      "bodyColumn": {
        "type": "string"
      }
    }
  }
}

```



```

    },
    "sqlQuery": {
      "type": "string",
      "not": {
        "pattern": ";+"
      }
    },
    },
    "timestampColumn": {
      "type": "string"
    },
    },
    "timestampFormat": {
      "type": "string"
    },
    },
    "timezone": {
      "type": "string"
    },
    },
    "changeDetectingColumns": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    },
    "allowedUsersColumn": {
      "type": "string"
    },
    },
    "allowedGroupsColumn": {
      "type": "string"
    },
    },
    "sourceURIColumn": {
      "type": "string"
    },
    },
    "serverlessAurora": {
      "type": "string",
      "enum": ["true", "false"]
    }
  },
  "required": ["primaryKey", "titleColumn", "bodyColumn", "sqlQuery"]
},
"type" : {
  "type" : "string",
  "pattern": "JDBC"
},
"syncMode": {
  "type": "string",

```

```

    "enum": [
      "FORCED_FULL_CRAWL",
      "FULL_CRAWL"
    ]
  },
  "secretArn": {
    "type": "string"
  }
},
"version": {
  "type": "string",
  "anyOf": [
    {
      "pattern": "1.0.0"
    }
  ]
},
"required": [
  "connectionConfiguration",
  "repositoryConfigurations",
  "syncMode",
  "additionalProperties",
  "secretArn",
  "type"
]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.
repositoryEndpointMetadata	Required configuration information for connecting your data source. <ul style="list-style-type: none"> dbType—The type of Java database you are using, whether <code>mysql</code>, <code>db2</code>, <code>postgresql</code>, <code>oracle</code>, or <code>sqlserver</code>. dbHost—The database host name.

Configuration	Description
	<ul style="list-style-type: none"> • dbPort—The database port. • dbInstance—The database instance.
repositoryConfigurations	<p>Configuration information for the content of the data source. For example, configuring specific types of content and field mappings. Specify the type of data source and the secret ARN.</p>
document	<p>A list of objects that map the attributes or field names of your database content to Amazon Q index field names. For more information, see Mapping data source fields.</p>
additionalProperties	<p>Additional configuration options for your content in your data source. Use to include or exclude specific content in your database data source.</p>
primaryKey	<p>Provide the primary key for the database table. This identifies a table within your database.</p>
titleColumn	<p>Provide the name of the document title column within your database table.</p>
bodyColumn	<p>Provide the name of the document title column within your database table.</p>

Configuration	Description
sqlQuery	Enter SQL query statements like SELECT and JOIN operations. SQL queries must be less than 1000 characters and not contain any semi-colons (;). Amazon Q will crawl all database content that matches your query. If a table name has special characters, put it in square brackets "[]" in the SQL query. For example: <code>select * from [my-database-table] .</code>
timestampColumn	Enter the name of the column which contains time stamps. Amazon Q uses time stamp information to detect changes in your content and sync only changed content.
timestampFormat	Enter the name of the column which contains time stamp formats to use to detect content changes and re-sync your content.
timezone	Enter the name of the column which contains time zones for the content to be crawled.
changeDetectingColumns	Enter the names of the columns that Amazon Q will use to detect content changes. Amazon Q will re-index content when there is a change in any of these columns
allowedUsersColumns	Enter the name of the column which contains User IDs to be allowed access to content.
allowedGroupsColumn	Enter the name of the column which contains User IDs to be allowed access to content.
sourceURIColumn	Enter the name of the column which contains Source URLs to be indexed.

Configuration	Description
isSslEnabled	true to add a path to an SSL certificate file stored in an Amazon S3 bucket.
type	The type of data source. Specify JDBC as your data source type.
syncMode	<p>Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose</p> <ul style="list-style-type: none">• <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index• <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index• <code>CHANGE_LOG</code> to incrementally crawl only new and modified content each time your data source syncs with your index.
secretArn	<p>The Amazon Resource Name (ARN) of a Secrets Manager secret that contains user name and password required to connect to your database. The secret must contain a JSON structure with the following keys:</p> <pre data-bbox="829 1503 1507 1696">{ "user name": "<i>database user name</i>", "password": "<i>password</i>" }</pre>
version	The version of the template that is currently supported.

How Amazon Q Business connector crawls Amazon RDS (Microsoft SQL Server) ACLs

When you connect a database data source to Amazon Q, Amazon Q crawls user and group information from a column in the source table. You specify this column in the console or using the configuration parameter as part of the `CreateDataSource` operation.

If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

A database data source has the following limitations:

- You can only specify an allow list for a database data source. You can't specify a deny list.
- You can only specify groups. You can't specify individual users for the allow list.
- The database column should be a string containing a semicolon delimited list of groups.

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business Amazon RDS (Microsoft SQL Server) data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a attribute mapping already available, or if you want to map additional document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q PostgreSQL connector supports the following field mappings:

Supported field mappings

- [Document](#)

Document

JDBC field name	Index field name	Description	Data type
jd_document_id	jd_document_id	Custom	String
jd_document_title	jd_document_title	Custom	String
jd_source_uri	_source_uri	Default	String

IAM role for Amazon Q Business Amazon RDS (Microsoft SQL Server) connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the `BatchPutDocument` and `BatchDeleteDocument` operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- Permission to access the SSL certificate stored in your Amazon S3 bucket.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AllowsAmazonQToGetS3Objects",
    "Action": [
      "s3:GetObject"
    ],
    "Resource": [
      "arn:aws:s3:::{{input_bucket_name}}/*"
    ],
    "Effect": "Allow",
    "Condition": {
      "StringEquals": {
        "aws:ResourceAccount": "{{account_id}}"
      }
    }
  }],
  {
    "Sid": "AllowsAmazonQToGetSecret",
    "Effect": "Allow",
    "Action": [
      "secretsmanager:GetSecretValue"
    ],
    "Resource": [
      "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
    ]
  }
}
```



```

    ],
  },
  {
    "Sid": "AllowsAmazonQToDecryptSecret",
    "Effect": "Allow",
    "Action": [
      "kms:Decrypt"
    ],
    "Resource": [
      "arn:aws:kms:{{region}}:{{account_id}}:key/[key_id]"
    ],
    "Condition": {
      "StringLike": {
        "kms:ViaService": [
          "secretsmanager.*.amazonaws.com"
        ]
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToIngestDocuments",
    "Effect": "Allow",
    "Action": [
      "qbusiness:BatchPutDocument",
      "qbusiness:BatchDeleteDocument"
    ],
    "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
    {{application_id}}/index/{{index_id}}"
  },
  {
    "Sid": "AllowsAmazonQToIngestPrincipalMapping",
    "Effect": "Allow",
    "Action": [
      "qbusiness:PutGroup",
      "qbusiness:CreateUser",
      "qbusiness>DeleteGroup",
      "qbusiness:UpdateUser",
      "qbusiness:ListGroups"
    ],
    "Resource": [
      "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
    {{application_id}}",
      "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
    {{application_id}}/index/{{index_id}}",

```

```

        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}/data-source/*"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateAndDeleteNI",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2>DeleteNetworkInterface"
    ],
    "Resource": [
      "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
      "arn:aws:ec2:{{region}}:{{account_id}}:security-group/
[[security_group]]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2>DeleteNetworkInterface"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:RequestTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
      },
      "ForAllValues:StringEquals": {
        "aws:TagKeys": [
          "AMAZON_Q"
        ]
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateTags",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateTags"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",

```

```

        "Condition": {
            "StringEquals": {
                "ec2:CreateAction": "CreateNetworkInterface"
            }
        },
        {
            "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
            "Effect": "Allow",
            "Action": [
                "ec2:CreateNetworkInterfacePermission"
            ],
            "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
            "Condition": {
                "StringLike": {
                    "aws:ResourceTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
                }
            }
        },
        {
            "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
            "Effect": "Allow",
            "Action": [
                "ec2:DescribeNetworkInterfaces",
                "ec2:DescribeAvailabilityZones",
                "ec2:DescribeNetworkInterfaceAttribute",
                "ec2:DescribeVpcs",
                "ec2:DescribeRegions",
                "ec2:DescribeNetworkInterfacePermissions",
                "ec2:DescribeSubnets"
            ],
            "Resource": "*"
        }
    ]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {

```

```

    "Sid": "AllowsAmazonQToAssumeRoleForServicePrincipal",
    "Effect": "Allow",
    "Principal": {
      "Service": "qbusiness.amazonaws.com"
    },
    "Action": "sts:AssumeRole",
    "Condition": {
      "StringEquals": {
        "aws:SourceAccount": "{{source_account}}"
      },
      "ArnLike": {
        "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
      }
    }
  }
]
}

```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Known limitations for the Amazon Q Business Amazon RDS (Microsoft SQL Server) connector

- Deleted database rows will not be tracked in when Amazon Q checks for updated content.
- The size of field names and values in a row of your database can't exceed 400KB.
- If you have a large amount of data in your database data source, and do not want Amazon Q to index all your database content after the first sync, you can choose to sync only new, modified, or deleted documents.

Connecting Amazon RDS (MySQL) to Amazon Q Business

Amazon RDS (MySQL) (Amazon Relational Database Service) is a web service that makes it easier to set up, operate, and scale a relational database in the AWS Cloud. You can connect your Amazon RDS (MySQL) instance to Amazon Q Business – using either the AWS Management Console, CLI, or the [CreateDataSource](#) API – and create an Amazon Q web experience.

The Amazon Q Aurora (MySQL) data source connector supports Amazon RDS MySql 5.6, 5.7, and 8.0.

⚠ Important

As a best practice, provide Amazon Q with read-only database credentials. Also, avoid adding tables with sensitive data or personal identifiable information (PII).

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).
- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [Amazon RDS \(MySQL\) connector overview](#)
- [Prerequisites for connecting Amazon Q Business to Amazon RDS \(MySQL\)](#)
- [Connecting Amazon Q Business to Amazon RDS \(MySQL\) using the console](#)
- [Connecting Amazon Q Business to Amazon RDS \(MySQL\) using APIs](#)
- [How Amazon Q Business connector crawls Amazon RDS \(MySQL\) ACLs](#)
- [Amazon Q Business Amazon RDS \(MySQL\) data source connector field mappings](#)
- [IAM role for Amazon Q Business Amazon RDS \(MySQL\) connector](#)
- [Known limitations for the Amazon Q Business Amazon RDS \(MySQL\) connector](#)

Amazon RDS (MySQL) connector overview

The following table gives an overview of the Amazon Q Business Amazon RDS (MySQL) connector and its supported features.

Category	Feature	Support
Security	Authentication type	Basic
	Authentication credentials	<ul style="list-style-type: none"> • Username of database user

Category	Feature	Support
		<ul style="list-style-type: none"> Password of database user
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Driver version	MySQL – 8.0.27
	Data source version	MySQL 5.6, 5.7, 8.0
	Identity crawling	No
	VPC	Yes
Crawl features	Custom metadata	Yes
	Entities	<p>Yes. The following entities are supported:</p> <ul style="list-style-type: none"> Document <div data-bbox="862 980 1510 1247" style="border: 1px solid #add8e6; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p>Note</p> <p>Each database row is considered an individual searchable Amazon Q document.</p> </div>
	Field mappings	Yes. Supports both default and custom field mappings. For more information, see Field mappings .
	Sync mode	Supports full and incremental sync.
	File types	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to Amazon RDS (MySQL)

Before you begin, make sure that you have completed the following prerequisites.

In Amazon RDS (MySQL), make sure you have:

- Noted your database user name and password.

Important

As a best practice, provide Amazon Q with read-only database credentials.

- Copied your database host URL, port, and instance. You can find this information on the Amazon RDS console.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your Amazon RDS (MySQL) authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Connecting Amazon Q Business to Amazon RDS (MySQL) using the console

The following procedure outlines how to connect Amazon Q Business to Amazon RDS (MySQL) using the AWS Management Console.

Connecting Amazon Q to Amazon RDS (MySQL)

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **Amazon RDS (MySQL)** page, enter the following information:
6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. In **Source**, enter the following information:
 - a. **Host** – Enter the database host URL, for example: `http://instance URL.region.rds.amazonaws.com`.
 - b. **Port** – Enter the database port, for example, 5432.
 - c. **Instance** – Enter the database instance, for example `postgres`.
 - d. **SSL certificate location** – Choose to enter the Amazon S3 path to your SSL certificate file.
8. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

 **Note**

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

9. In **Authentication** – Enter the following information for your **AWS Secrets Manager secret**.
 - a. **Secret name** – A name for your secret.
 - b. For **Database user name**, and **Password** – Enter the authentication credential values you copied from your database.
 - c. Choose **Save**.
10. **Configure VPC and security group – optional** – Choose whether you want to use a VPC. If you do, enter the following information:
 - a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.

- b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

11. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

12. In **Sync scope**, enter the following information:

- **SQL query** – Enter SQL query statements like SELECT and JOIN operations. SQL queries must be less than 1000 characters and not contain any semi-colons (;). Amazon Q will crawl all database content that matches your query.
- **Primary key column** – Provide the primary key for the database table. This identifies a table within your database.
- **Title column** – Provide the name of the document title column within your database table.
- **Body column** – Provide the name of the document body column within your database table.

13. In **Additional configuration – optional** – Configure the following settings:

- **Change-detecting columns** – Enter the names of the columns that Amazon Q will use to detect content changes. Amazon Q will re-index content when there is a change in any of these columns.
- **Users' IDs column** – Enter the name of the column which contains User IDs to be allowed access to content.
- **Groups column** – Enter the name of the column that contains groups to be allowed access to content.
- **Source URLs column** – Enter the name of the column which contains Source URLs to be indexed.
- **Time stamps column** – Enter the name of the column which contains time stamps. Amazon Q uses time stamp information to detect changes in your content and sync only changed content.
- **Time zones column** – Enter the name of the column which contains time zones for the content to be crawled.

- **Time stamps format** – Enter the name of the column which contains time stamp formats to use to detect content changes and re-sync your content.
14. In **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.
- **Full sync** – Sync all content regardless of the previous sync status.
 - **New or modified content sync** – Sync only new and modified documents.
 - **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.

For more details, see [Sync mode](#).

15. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
16. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
17. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:
- a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
 - b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

 **Note**

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

18. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

19. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

Note

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to Amazon RDS (MySQL) using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the configuration parameter to provide a JSON schema with all other configuration information specific to your data source connector.

Amazon RDS (MySQL) JSON schema

The following is the Amazon RDS (MySQL) JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
      "type": "object",
      "properties": {
        "repositoryEndpointMetadata": {
          "type": "object",
          "properties": {
            "dbType": {
              "type": "string",
              "enum": [
                "mysql",
                "db2",
                "postgresql",
                "oracle",
                "sqlserver"
              ]
            }
          }
        }
      }
    }
  }
}
```

```

    },
    "dbHost": {
      "type": "string"
    },
    "dbPort": {
      "type": "string"
    },
    "dbInstance": {
      "type": "string"
    }
  },
  "required": [
    "dbType",
    "dbHost",
    "dbPort",
    "dbInstance"
  ]
}
},
"required": [
  "repositoryEndpointMetadata"
]
},
"repositoryConfigurations": {
  "type": "object",
  "properties": {
    "document": {
      "type": "object",
      "properties": {
        "fieldMappings": {
          "type": "array",
          "items": [
            {
              "type": "object",
              "properties": {
                "indexFieldName": {
                  "type": "string"
                },
                "indexFieldType": {
                  "type": "string"
                },
                "dataSourceFieldName": {
                  "type": "string"
                }
              }
            }
          ]
        }
      }
    }
  }
}

```

```

        },
        "required": [
            "indexFieldName",
            "indexFieldType",
            "dataSourceFieldName"
        ]
    }
]
}
},
"required": [
    "fieldMappings"
]
}
},
"required": [
]
},
"additionalProperties": {
    "type": "object",
    "properties": {
        "primaryKey": {
            "type": "string"
        },
        "titleColumn": {
            "type": "string"
        },
        "bodyColumn": {
            "type": "string"
        },
        "sqlQuery": {
            "type": "string",
            "not": {
                "pattern": ";+"
            }
        },
        "timestampColumn": {
            "type": "string"
        },
        "timestampFormat": {
            "type": "string"
        },
        "timezone": {
            "type": "string"
        }
    }
}

```

```
    },
    "changeDetectingColumns": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    "allowedUsersColumn": {
      "type": "string"
    },
    "allowedGroupsColumn": {
      "type": "string"
    },
    "sourceURIColumn": {
      "type": "string"
    },
    "serverlessAurora": {
      "type": "string",
      "enum": ["true", "false"]
    }
  },
  "required": ["primaryKey", "titleColumn", "bodyColumn", "sqlQuery"]
},
"type" : {
  "type" : "string",
  "pattern": "JDBC"
},
"syncMode": {
  "type": "string",
  "enum": [
    "FORCED_FULL_CRAWL",
    "FULL_CRAWL"
  ]
},
"secretArn": {
  "type": "string"
}
},
"version": {
  "type": "string",
  "anyOf": [
    {
      "pattern": "1.0.0"
    }
  ]
}
```

```

    ]
  },
  "required": [
    "connectionConfiguration",
    "repositoryConfigurations",
    "syncMode",
    "additionalProperties",
    "secretArn",
    "type"
  ]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.
repositoryEndpointMetadata	<p>Required configuration information for connecting your data source.</p> <ul style="list-style-type: none"> dbType—The type of Java database you are using, whether <code>mysql</code>, <code>db2</code>, <code>postgresql</code>, <code>oracle</code>, or <code>sqlserver</code>. dbHost—The database host name. dbPort—The database port. dbInstance—The database instance.
repositoryConfigurations	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings. Specify the type of data source and the secret ARN.
document	A list of objects that map the attributes or field names of your database content to Amazon Q index field names. For more information, see Mapping data source fields .

Configuration	Description
additionalProperties	Additional configuration options for your content in your data source. Use to include or exclude specific content in your database data source.
primaryKey	Provide the primary key for the database table. This identifies a table within your database.
titleColumn	Provide the name of the document title column within your database table.
bodyColumn	Provide the name of the document title column within your database table.
sqlQuery	Enter SQL query statements like SELECT and JOIN operations. SQL queries must be less than 1000 characters and not contain any semi-colons (;). Amazon Q will crawl all database content that matches your query.
timestampColumn	Enter the name of the column which contains time stamps. Amazon Q uses time stamp information to detect changes in your content and sync only changed content.
timestampFormat	Enter the name of the column which contains time stamp formats to use to detect content changes and re-sync your content.
timezone	Enter the name of the column which contains time zones for the content to be crawled.
changeDetectingColumns	Enter the names of the columns that Amazon Q will use to detect content changes. Amazon Q will re-index content when there is a change in any of these columns

Configuration	Description
allowedUsersColumns	Enter the name of the column which contains User IDs to be allowed access to content.
allowedGroupsColumn	Enter the name of the column which contains User IDs to be allowed access to content.
sourceURIColumn	Enter the name of the column which contains Source URLs to be indexed.
isSslEnabled	<code>true</code> to add a path to an SSL certificate file stored in an Amazon S3 bucket.
type	The type of data source. Specify JDBC as your data source type.
syncMode	Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose <ul style="list-style-type: none"><code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index<code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index<code>CHANGE_LOG</code> to incrementally crawl only new and modified content each time your data source syncs with your index.

Configuration	Description
secretArn	The Amazon Resource Name (ARN) of a Secrets Manager secret that contains user name and password required to connect to your database. The secret must contain a JSON structure with the following keys: <pre>{ "user name": "<i>database user name</i>", "password": "<i>password</i>" }</pre>
version	The version of the template that is currently supported.

How Amazon Q Business connector crawls Amazon RDS (MySQL) ACLs

When you connect a database data source to Amazon Q, Amazon Q crawls user and group information from a column in the source table. You specify this column in the console or using the configuration parameter as part of the `CreateDataSource` operation.

If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

A database data source has the following limitations:

- You can only specify an allow list for a database data source. You can't specify a deny list.
- You can only specify groups. You can't specify individual users for the allow list.
- The database column should be a string containing a semicolon delimited list of groups.

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business Amazon RDS (MySQL) data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a attribute mapping already available, or if you want to map additional document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q PostgreSQL connector supports the following field mappings:

Supported field mappings

- [Document](#)

Document

JDBC field name	Index field name	Description	Data type
jd_document_id	jd_document_id	Custom	String
jd_document_title	jd_document_title	Custom	String
jd_source_uri	_source_uri	Default	String

IAM role for Amazon Q Business Amazon RDS (MySQL) connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the `BatchPutDocument` and `BatchDeleteDocument` operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- Permission to access the SSL certificate stored in your Amazon S3 bucket.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AllowsAmazonQToGetS3Objects",
    "Action": [
      "s3:GetObject"
    ],
  },
```

```

    "Resource": [
      "arn:aws:s3:::{{input_bucket_name}}/*"
    ],
    "Effect": "Allow",
    "Condition": {
      "StringEquals": {
        "aws:ResourceAccount": "{{account_id}}"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToGetSecret",
    "Effect": "Allow",
    "Action": [
      "secretsmanager:GetSecretValue"
    ],
    "Resource": [
      "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToDecryptSecret",
    "Effect": "Allow",
    "Action": [
      "kms:Decrypt"
    ],
    "Resource": [
      "arn:aws:kms:{{region}}:{{account_id}}:key/[key_id]"
    ],
    "Condition": {
      "StringLike": {
        "kms:ViaService": [
          "secretsmanager.*.amazonaws.com"
        ]
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToIngestDocuments",
    "Effect": "Allow",
    "Action": [
      "qbusiness:BatchPutDocument",
      "qbusiness:BatchDeleteDocument"
    ],
  },

```

```

    "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
  },
  {
    "Sid": "AllowsAmazonQToIngestPrincipalMapping",
    "Effect": "Allow",
    "Action": [
      "qbusiness:PutGroup",
      "qbusiness:CreateUser",
      "qbusiness>DeleteGroup",
      "qbusiness:UpdateUser",
      "qbusiness:ListGroups"
    ],
    "Resource": [
      "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}",
      "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}",
      "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}/data-source/*"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateAndDeleteNI",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2>DeleteNetworkInterface"
    ],
    "Resource": [
      "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
      "arn:aws:ec2:{{region}}:{{account_id}}:security-group/
[{{security_group}}]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2>DeleteNetworkInterface"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {

```

```

        "StringLike": {
            "aws:RequestTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
        },
        "ForAllValues:StringEquals": {
            "aws:TagKeys": [
                "AMAZON_Q"
            ]
        }
    },
    {
        "Sid": "AllowsAmazonQToCreateTags",
        "Effect": "Allow",
        "Action": [
            "ec2:CreateTags"
        ],
        "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
        "Condition": {
            "StringEquals": {
                "ec2:CreateAction": "CreateNetworkInterface"
            }
        }
    },
    {
        "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
        "Effect": "Allow",
        "Action": [
            "ec2:CreateNetworkInterfacePermission"
        ],
        "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
        "Condition": {
            "StringLike": {
                "aws:ResourceTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
            }
        }
    },
    {
        "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
        "Effect": "Allow",
        "Action": [
            "ec2:DescribeNetworkInterfaces",
            "ec2:DescribeAvailabilityZones",

```

```

        "ec2:DescribeNetworkInterfaceAttribute",
        "ec2:DescribeVpcs",
        "ec2:DescribeRegions",
        "ec2:DescribeNetworkInterfacePermissions",
        "ec2:DescribeSubnets"
    ],
    "Resource": "*"
}
]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToAssumeRoleForServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "{{source_account}}"
        },
        "ArnLike": {
          "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
        }
      }
    }
  ]
}

```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Known limitations for the Amazon Q Business Amazon RDS (MySQL) connector

- Deleted database rows will not be tracked in when Amazon Q checks for updated content.

- The size of field names and values in a row of your database can't exceed 400KB.
- If you have a large amount of data in your database data source, and do not want Amazon Q to index all your database content after the first sync, you can choose to sync only new, modified, or deleted documents.

Connecting Amazon RDS (Oracle) to Amazon Q Business

Amazon RDS (Oracle) (Amazon Relational Database Service) is a web service that makes it easier to set up, operate, and scale a relational database in the AWS Cloud. You can connect your Amazon RDS (Oracle) instance to Amazon Q Business – using either the AWS Management Console, CLI, or the [CreateDataSource](#) API – and create an Amazon Q web experience.

The Amazon Q Aurora (MySQL) data source connector supports Amazon RDS Oracle Database 21c, Oracle Database 19c, Oracle Database 12c.

Important

As a best practice, provide Amazon Q with read-only database credentials. Also, avoid adding tables with sensitive data or personal identifiable information (PII).

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).
- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [Amazon RDS \(Oracle\) connector overview](#)
- [Prerequisites for connecting Amazon Q Business to Amazon RDS \(Oracle\)](#)
- [Connecting Amazon Q Business to Amazon RDS \(Oracle\) using the console](#)
- [Connecting Amazon Q Business to Amazon RDS \(Oracle\) using APIs](#)
- [How Amazon Q Business connector crawls Amazon RDS \(Oracle\) ACLs](#)

- [Amazon Q Business Amazon RDS \(Oracle\) data source connector field mappings](#)
- [IAM role for Amazon Q Business Amazon RDS \(Oracle\) connector](#)
- [Known limitations for the Amazon Q Business Amazon RDS \(Oracle\) connector](#)

Amazon RDS (Oracle) connector overview

The following table gives an overview of the Amazon Q Business Amazon RDS (Oracle) connector and its supported features.

Category	Feature	Support
Security	Authentication type	Basic
	Authentication credentials	<ul style="list-style-type: none"> • Username of database user • Password of database user
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Driver version	Oracle – 21.1.0.0
	Data source version	Oracle Database 12c, 19c, 21c
	Identity crawling	No
	VPC	Yes
Crawl features	Custom metadata	Yes
	Entities	<p>Yes. The following entities are supported:</p> <ul style="list-style-type: none"> • Document <div style="border: 1px solid #00a0e3; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p>Note</p> <p>Each database row is considered an individual searchable Amazon Q document.</p> </div>

Category	Feature	Support
	Field mappings	Yes. Supports both default and custom field mappings. For more information, see Field mappings .
	Sync mode	Supports full and incremental sync.
	File types	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to Amazon RDS (Oracle)

Before you begin, make sure that you have completed the following prerequisites.

In Amazon RDS (Oracle), make sure you have:

- Noted your database user name and password.

Important

As a best practice, provide Amazon Q with read-only database credentials.

- Copied your database host URL, port, and instance.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your Amazon RDS (Oracle) authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Connecting Amazon Q Business to Amazon RDS (Oracle) using the console

The following procedure outlines how to connect Amazon Q Business to Amazon RDS (Oracle) using the AWS Management Console.

Connecting Amazon Q to Amazon RDS (Oracle)

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **Amazon RDS (Oracle)** page, enter the following information:
6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. **Source**, enter the following information:
 - a. **Host** – Enter the database host name.
 - b. **Port** – Enter the database port.
 - c. **Instance** – Enter the database instance.
 - d. **SSL certificate location** – Choose to enter the Amazon S3 path to your SSL certificate file.
8. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

Note

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

9. In **Authentication** – Enter the following information for your **AWS Secrets Manager secret**.
 - a. **Secret name** – A name for your secret.
 - b. For **Database user name**, and **Password** – Enter the authentication credential values you copied from your database.
 - c. Choose **Save**.
10. **Configure VPC and security group – optional** – Choose whether you want to use a VPC. If you do, enter the following information:
 - a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
 - b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

11. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

12. In **Sync scope**, enter the following information:
 - **SQL query** – Enter SQL query statements like SELECT and JOIN operations. SQL queries must be less than 1000 characters and not contain any semi-colons (;). Amazon Q will crawl all database content that matches your query.
 - **Primary key column** – Provide the primary key for the database table. This identifies a table within your database.
 - **Title column** – Provide the name of the document title column within your database table.
 - **Body column** – Provide the name of the document body column within your database table.
13. In **Additional configuration – optional** – Configure the following settings:
 - **Change-detecting columns** – Enter the names of the columns that Amazon Q will use to detect content changes. Amazon Q will re-index content when there is a change in any of these columns.
 - **Users' IDs column** – Enter the name of the column which contains User IDs to be allowed access to content.

- **Groups column** – Enter the name of the column that contains groups to be allowed access to content.
 - **Source URLs column** – Enter the name of the column which contains Source URLs to be indexed.
 - **Time stamps column** – Enter the name of the column which contains time stamps. Amazon Q uses time stamp information to detect changes in your content and sync only changed content.
 - **Time zones column** – Enter the name of the column which contains time zones for the content to be crawled.
 - **Time stamps format** – Enter the name of the column which contains time stamp formats to use to detect content changes and re-sync your content.
14. In **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.
- **Full sync** – Sync all content regardless of the previous sync status.
 - **New or modified content sync** – Sync only new and modified documents.
 - **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.

For more details, see [Sync mode](#).

15. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
16. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
17. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:
- a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
 - b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

Note

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

18. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

19. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

Note

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to Amazon RDS (Oracle) using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the `configuration` parameter to provide a JSON schema with all other configuration information specific to your data source connector.

Amazon RDS (Oracle) JSON schema

The following is the Amazon RDS (Oracle) JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
```

```

"type": "object",
"properties": {
  "repositoryEndpointMetadata": {
    "type": "object",
    "properties": {
      "dbType": {
        "type": "string",
        "enum": [
          "mysql",
          "db2",
          "postgresql",
          "oracle",
          "sqlserver"
        ]
      },
      "dbHost": {
        "type": "string"
      },
      "dbPort": {
        "type": "string"
      },
      "dbInstance": {
        "type": "string"
      }
    },
    "required": [
      "dbType",
      "dbHost",
      "dbPort",
      "dbInstance"
    ]
  },
  "required": [
    "repositoryEndpointMetadata"
  ]
},
"repositoryConfigurations": {
  "type": "object",
  "properties": {
    "document": {
      "type": "object",
      "properties": {
        "fieldMappings": {

```



```
    "type": "array",
    "items": [
      {
        "type": "object",
        "properties": {
          "indexFieldName": {
            "type": "string"
          },
          "indexFieldType": {
            "type": "string"
          },
          "dataSourceFieldName": {
            "type": "string"
          }
        },
        "required": [
          "indexFieldName",
          "indexFieldType",
          "dataSourceFieldName"
        ]
      }
    ]
  },
  "required": [
    "fieldMappings"
  ]
},
"required": [
]
},
"additionalProperties": {
  "type": "object",
  "properties": {
    "primaryKey": {
      "type": "string"
    },
    "titleColumn": {
      "type": "string"
    },
    "bodyColumn": {
      "type": "string"
    }
  },
}
```

```

    "sqlQuery": {
      "type": "string",
      "not": {
        "pattern": ";+"
      }
    },
    "timestampColumn": {
      "type": "string"
    },
    "timestampFormat": {
      "type": "string"
    },
    "timezone": {
      "type": "string"
    },
    "changeDetectingColumns": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    "allowedUsersColumn": {
      "type": "string"
    },
    "allowedGroupsColumn": {
      "type": "string"
    },
    "sourceURIColumn": {
      "type": "string"
    },
    "serverlessAurora": {
      "type": "string",
      "enum": ["true", "false"]
    }
  },
  "required": ["primaryKey", "titleColumn", "bodyColumn", "sqlQuery"]
},
"type" : {
  "type" : "string",
  "pattern": "JDBC"
},
"syncMode": {
  "type": "string",
  "enum": [

```

```

        "FORCED_FULL_CRAWL",
        "FULL_CRAWL"
    ]
},
"secretArn": {
    "type": "string"
}
},
"version": {
    "type": "string",
    "anyOf": [
        {
            "pattern": "1.0.0"
        }
    ]
}
},
"required": [
    "connectionConfiguration",
    "repositoryConfigurations",
    "syncMode",
    "additionalProperties",
    "secretArn",
    "type"
]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.
repositoryEndpointMetadata	<p>Required configuration information for connecting your data source.</p> <ul style="list-style-type: none"> dbType—The type of Java database you are using, whether <code>mysql</code>, <code>db2</code>, <code>postgresql</code>, <code>oracle</code>, or <code>sqlserver</code>. dbHost—The database host name. dbPort—The database port.

Configuration	Description
	<ul style="list-style-type: none"> • <code>dbInstance</code>—The database instance.
<code>repositoryConfigurations</code>	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings. Specify the type of data source and the secret ARN.
<code>document</code>	A list of objects that map the attributes or field names of your database content to Amazon Q index field names. For more information, see Mapping data source fields .
<code>additionalProperties</code>	Additional configuration options for your content in your data source. Use to include or exclude specific content in your database data source.
<code>primaryKey</code>	Provide the primary key for the database table. This identifies a table within your database.
<code>titleColumn</code>	Provide the name of the document title column within your database table.
<code>bodyColumn</code>	Provide the name of the document title column within your database table.
<code>sqlQuery</code>	Enter SQL query statements like <code>SELECT</code> and <code>JOIN</code> operations. SQL queries must be less than 1000 characters and not contain any semi-colons (;). Amazon Q will crawl all database content that matches your query.

Configuration	Description
timestampColumn	Enter the name of the column which contains time stamps. Amazon Q uses time stamp information to detect changes in your content and sync only changed content.
timestampFormat	Enter the name of the column which contains time stamp formats to use to detect content changes and re-sync your content.
timezone	Enter the name of the column which contains time zones for the content to be crawled.
changeDetectingColumns	Enter the names of the columns that Amazon Q will use to detect content changes. Amazon Q will re-index content when there is a change in any of these columns
allowedUsersColumns	Enter the name of the column which contains User IDs to be allowed access to content.
allowedGroupsColumn	Enter the name of the column which contains User IDs to be allowed access to content.
sourceURIColumn	Enter the name of the column which contains Source URLs to be indexed.
isSslEnabled	<code>true</code> to add a path to an SSL certificate file stored in an Amazon S3 bucket.
type	The type of data source. Specify JDBC as your data source type.

Configuration	Description
syncMode	<p>Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose</p> <ul style="list-style-type: none"> • <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index • <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index • <code>CHANGE_LOG</code> to incrementally crawl only new and modified content each time your data source syncs with your index.
secretArn	<p>The Amazon Resource Name (ARN) of a Secrets Manager secret that contains user name and password required to connect to your database. The secret must contain a JSON structure with the following keys:</p> <pre data-bbox="829 1247 1507 1440"> { "user name": "<i>database user name</i>", "password": "<i>password</i>" } </pre>
version	<p>The version of the template that is currently supported.</p>

How Amazon Q Business connector crawls Amazon RDS (Oracle) ACLs

When you connect a database data source to Amazon Q, Amazon Q crawls user and group information from a column in the source table. You specify this column in the console or using the configuration parameter as part of the `CreateDataSource` operation.

If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

A database data source has the following limitations:

- You can only specify an allow list for a database data source. You can't specify a deny list.
- You can only specify groups. You can't specify individual users for the allow list.
- The database column should be a string containing a semicolon delimited list of groups.

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business Amazon RDS (Oracle) data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a attribute mapping already available, or if you want to map additional document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

⚠ Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q PostgreSQL connector supports the following field mappings:

Supported field mappings

- [Document](#)

Document

JDBC field name	Index field name	Description	Data type
jd_document_id	jd_document_id	Custom	String
jd_document_title	jd_document_title	Custom	String
jd_source_uri	_source_uri	Default	String

IAM role for Amazon Q Business Amazon RDS (Oracle) connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the BatchPutDocument and BatchDeleteDocument operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.

- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- Permission to access the SSL certificate stored in your Amazon S3 bucket.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AllowsAmazonQToGetS3Objects",
    "Action": [
      "s3:GetObject"
    ],
    "Resource": [
      "arn:aws:s3:::{{input_bucket_name}}/*"
    ],
    "Effect": "Allow",
    "Condition": {
      "StringEquals": {
        "aws:ResourceAccount": "{{account_id}}"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToGetSecret",
    "Effect": "Allow",
    "Action": [
      "secretsmanager:GetSecretValue"
    ],
    "Resource": [
      "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToDecryptSecret",
    "Effect": "Allow",
    "Action": [
      "kms:Decrypt"
    ],
    "Resource": [
      "arn:aws:kms:{{region}}:{{account_id}}:key/[key_id]"
    ]
  },

```

```

        "Condition": {
            "StringLike": {
                "kms:ViaService": [
                    "secretsmanager.*.amazonaws.com"
                ]
            }
        },
        {
            "Sid": "AllowsAmazonQToIngestDocuments",
            "Effect": "Allow",
            "Action": [
                "qbusiness:BatchPutDocument",
                "qbusiness:BatchDeleteDocument"
            ],
            "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
        },
        {
            "Sid": "AllowsAmazonQToIngestPrincipalMapping",
            "Effect": "Allow",
            "Action": [
                "qbusiness:PutGroup",
                "qbusiness:CreateUser",
                "qbusiness>DeleteGroup",
                "qbusiness:UpdateUser",
                "qbusiness:ListGroup"
            ],
            "Resource": [
                "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}",
                "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}",
                "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}/data-source/*"
            ]
        },
        {
            "Sid": "AllowsAmazonQToCreateAndDeleteNI",
            "Effect": "Allow",
            "Action": [
                "ec2:CreateNetworkInterface",
                "ec2>DeleteNetworkInterface"
            ],

```

```

        "Resource": [
            "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
            "arn:aws:ec2:{{region}}:{{account_id}}:security-group/
[[security_group]]"
        ]
    },
    {
        "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
        "Effect": "Allow",
        "Action": [
            "ec2:CreateNetworkInterface",
            "ec2>DeleteNetworkInterface"
        ],
        "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
        "Condition": {
            "StringLike": {
                "aws:RequestTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
            },
            "ForAllValues:StringEquals": {
                "aws:TagKeys": [
                    "AMAZON_Q"
                ]
            }
        }
    },
    {
        "Sid": "AllowsAmazonQToCreateTags",
        "Effect": "Allow",
        "Action": [
            "ec2:CreateTags"
        ],
        "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
        "Condition": {
            "StringEquals": {
                "ec2:CreateAction": "CreateNetworkInterface"
            }
        }
    },
    {
        "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
        "Effect": "Allow",
        "Action": [
            "ec2:CreateNetworkInterfacePermission"
        ]
    }
}

```

```

    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:ResourceTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
    "Effect": "Allow",
    "Action": [
      "ec2:DescribeNetworkInterfaces",
      "ec2:DescribeAvailabilityZones",
      "ec2:DescribeNetworkInterfaceAttribute",
      "ec2:DescribeVpcs",
      "ec2:DescribeRegions",
      "ec2:DescribeNetworkInterfacePermissions",
      "ec2:DescribeSubnets"
    ],
    "Resource": "*"
  }
]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToAssumeRoleForServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "{{source_account}}"
        },
        "ArnLike": {

```

```
        "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
    }
}
]
}
```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Known limitations for the Amazon Q Business Amazon RDS (Oracle) connector

- Deleted database rows will not be tracked in when Amazon Q checks for updated content.
- The size of field names and values in a row of your database can't exceed 400KB.
- If you have a large amount of data in your database data source, and do not want Amazon Q to index all your database content after the first sync, you can choose to sync only new, modified, or deleted documents.

Connecting Amazon RDS (PostgreSQL) to Amazon Q Business

Amazon RDS (PostgreSQL) is a web service that makes it easier to set up, operate, and scale a relational database in the AWS Cloud. If you are a AWS user, you can use Amazon Q Business to index your Amazon RDS (PostgreSQL) data source.

The Amazon Q Amazon RDS (PostgreSQL) data source connector supports PostgreSQL 9.6.

You can connect your Amazon RDS (PostgreSQL) instance to Amazon Q Business—using either the AWS Management Console, CLI, or the [CreateDataSource](#) API—and create an Amazon Q web experience.

Important

As a best practice, provide Amazon Q with read-only database credentials. Also, avoid adding tables with sensitive data or personal identifiable information (PII).

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).
- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [Amazon RDS \(PostgreSQL\) connector overview](#)
- [Prerequisites for connecting Amazon Q Business to Amazon RDS \(PostgreSQL\)](#)
- [Connecting Amazon Q Business to Amazon RDS \(PostgreSQL\) using the console](#)
- [Connecting Amazon Q Business to Amazon RDS \(PostgreSQL\) using APIs](#)
- [How Amazon Q Business connector crawls Amazon RDS \(PostgreSQL\) ACLs](#)
- [Amazon Q Business Amazon RDS \(PostgreSQL\) data source connector field mappings](#)
- [IAM role for Amazon Q Business Amazon RDS \(PostgreSQL\) connector](#)
- [Known limitations for the Amazon Q Business Amazon RDS \(PostgreSQL\) connector](#)

Amazon RDS (PostgreSQL) connector overview

The following table gives an overview of the Amazon Q Business Amazon RDS (PostgreSQL) connector and its supported features.

Category	Feature	Support
Security	Authentication type	Basic
	Authentication credentials	<ul style="list-style-type: none"> • Username of database user • Password of database user
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Driver version	PostgreSQL – 42.3.2
	Data source version	PostgreSQL 9.6

Category	Feature	Support
	Identity crawling	No
	VPC	Yes
Crawl features	Custom metadata	Yes
	Entities	Yes. The following entities are supported: <ul style="list-style-type: none"> Document <div data-bbox="862 625 1507 892" style="border: 1px solid #add8e6; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p>Note</p> <p>Each database row is considered an individual searchable Amazon Q document.</p> </div>
	Field mappings	Yes. Supports both default and custom field mappings. For more information, see Field mappings .
	Sync mode	Supports full and incremental sync.
	File types	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to Amazon RDS (PostgreSQL)

Before you begin, make sure that you have completed the following prerequisites.

In Amazon RDS (PostgreSQL), make sure you have:

- Noted your database user name and password.

Important

As a best practice, provide Amazon Q with read-only database credentials.

- Copied your database host URL, port, and instance. You can find this information on the Amazon RDS console.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your Amazon RDS (PostgreSQL) authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Connecting Amazon Q Business to Amazon RDS (PostgreSQL) using the console

The following procedure outlines how to connect Amazon Q Business to Amazon RDS (PostgreSQL) using the AWS Management Console.

Connecting Amazon Q to Amazon RDS (PostgreSQL)

1. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

2. In **Source**, enter the following information:

- a. **Host** – Enter the database host URL, for example: `http://instance URL.region.rds.amazonaws.com`.

- b. **Port** – Enter the database port, for example, 5432.

- c. **Instance** – Enter the database instance, for example `postgres`.

- d. **SSL certificate location** – Choose to enter the Amazon S3 path to your SSL certificate file.

3. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

 **Note**

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

4. In **Authentication** – Enter the following information for your **AWS Secrets Manager secret**.
 - a. **Secret name** – A name for your secret.
 - b. For **Database user name**, and **Password** – Enter the authentication credential values you copied from your database.
 - c. Choose **Save**.
5. **Configure VPC and security group – optional** – Choose whether you want to use a VPC. If you do, enter the following information:
 - a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
 - b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

6. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

7. In **Sync scope**, enter the following information:
 - **SQL query** – Enter SQL query statements like SELECT and JOIN operations. SQL queries must be less than 1000 characters and not contain any semi-colons (;). Amazon Q will crawl all database content that matches your query.
 - **Primary key column** – Provide the primary key for the database table. This identifies a table within your database.

- **Title column** – Provide the name of the document title column within your database table.
 - **Body column** – Provide the name of the document body column within your database table.
8. In **Additional configuration** – *optional* – Configure the following settings:
- **Change-detecting columns** – Enter the names of the columns that Amazon Q will use to detect content changes. Amazon Q will re-index content when there is a change in any of these columns.
 - **Users' IDs column** – Enter the name of the column which contains User IDs to be allowed access to content.
 - **Groups column** – Enter the name of the column that contains groups to be allowed access to content.
 - **Source URLs column** – Enter the name of the column which contains Source URLs to be indexed.
 - **Time stamps column** – Enter the name of the column which contains time stamps. Amazon Q uses time stamp information to detect changes in your content and sync only changed content.
 - **Time zones column** – Enter the name of the column which contains time zones for the content to be crawled.
 - **Time stamps format** – Enter the name of the column which contains time stamp formats to use to detect content changes and re-sync your content.
9. In **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.
- **Full sync** – Sync all content regardless of the previous sync status.
 - **New or modified content sync** – Sync only new and modified documents.
 - **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.

For more details, see [Sync mode](#).

10. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
11. **Tags** - *optional* – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.

12. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:
 - a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
 - b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

 **Note**

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

13. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

14. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

 **Note**

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to Amazon RDS (PostgreSQL) using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the `configuration` parameter to provide a JSON schema with all other configuration information specific to your data source connector.

Amazon RDS (PostgreSQL) JSON schema

The following is the Amazon RDS (PostgreSQL) JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
      "type": "object",
      "properties": {
        "repositoryEndpointMetadata": {
          "type": "object",
          "properties": {
            "dbType": {
              "type": "string",
              "enum": [
                "mysql",
                "db2",
                "postgresql",
                "oracle",
                "sqlserver"
              ]
            },
            "dbHost": {
              "type": "string"
            },
            "dbPort": {
              "type": "string"
            },
            "dbInstance": {
              "type": "string"
            }
          },
          "required": [
            "dbType",
            "dbHost",
            "dbPort",
            "dbInstance"
          ]
        }
      },
      "required": [
        "repositoryEndpointMetadata"
      ]
    }
  }
}
```

```

]
},
"repositoryConfigurations": {
  "type": "object",
  "properties": {
    "document": {
      "type": "object",
      "properties": {
        "fieldMappings": {
          "type": "array",
          "items": [
            {
              "type": "object",
              "properties": {
                "indexFieldName": {
                  "type": "string"
                },
                "indexFieldType": {
                  "type": "string"
                },
                "dataSourceFieldName": {
                  "type": "string"
                }
              },
              "required": [
                "indexFieldName",
                "indexFieldType",
                "dataSourceFieldName"
              ]
            }
          ]
        }
      },
      "required": [
        "fieldMappings"
      ]
    }
  },
  "required": [
  ]
},
"additionalProperties": {
  "type": "object",
  "properties": {

```

```
"primaryKey": {
  "type": "string"
},
"titleColumn": {
  "type": "string"
},
"bodyColumn": {
  "type": "string"
},
"sqlQuery": {
  "type": "string",
  "not": {
    "pattern": ";+"
  }
},
"timestampColumn": {
  "type": "string"
},
"timestampFormat": {
  "type": "string"
},
"timezone": {
  "type": "string"
},
"changeDetectingColumns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"allowedUsersColumn": {
  "type": "string"
},
"allowedGroupsColumn": {
  "type": "string"
},
"sourceURIColumn": {
  "type": "string"
},
"serverlessAurora": {
  "type": "string",
  "enum": ["true", "false"]
}
},
```

```

    "required": ["primaryKey", "titleColumn", "bodyColumn", "sqlQuery"]
  },
  "type" : {
    "type" : "string",
    "pattern": "JDBC"
  },
  "syncMode": {
    "type": "string",
    "enum": [
      "FORCED_FULL_CRAWL",
      "FULL_CRAWL"
    ]
  },
  "secretArn": {
    "type": "string"
  }
},
"version": {
  "type": "string",
  "anyOf": [
    {
      "pattern": "1.0.0"
    }
  ]
},
"required": [
  "connectionConfiguration",
  "repositoryConfigurations",
  "syncMode",
  "additionalProperties",
  "secretArn",
  "type"
]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.

Configuration	Description
repositoryEndpointMetadata	<p>Required configuration information for connecting your data source.</p> <ul style="list-style-type: none"> • dbType—The type of Java database you are using, whether <code>mysql</code>, <code>db2</code>, <code>postgresql</code>, <code>oracle</code>, or <code>sqlserver</code>. • dbHost—The database host name. • dbPort—The database port. • dbInstance—The database instance.
repositoryConfigurations	<p>Configuration information for the content of the data source. For example, configuring specific types of content and field mappings. Specify the type of data source and the secret ARN.</p>
document	<p>A list of objects that map the attributes or field names of your database content to Amazon Q index field names. For more information, see Mapping data source fields.</p>
additionalProperties	<p>Additional configuration options for your content in your data source. Use to include or exclude specific content in your database data source.</p>
primaryKey	<p>Provide the primary key for the database table. This identifies a table within your database.</p>
titleColumn	<p>Provide the name of the document title column within your database table.</p>
bodyColumn	<p>Provide the name of the document title column within your database table.</p>

Configuration	Description
sqlQuery	Enter SQL query statements like SELECT and JOIN operations. SQL queries must be less than 1000 characters and not contain any semi-colons (;). Amazon Q will crawl all database content that matches your query.
timestampColumn	Enter the name of the column which contains time stamps. Amazon Q uses time stamp information to detect changes in your content and sync only changed content.
timestampFormat	Enter the name of the column which contains time stamp formats to use to detect content changes and re-sync your content.
timezone	Enter the name of the column which contains time zones for the content to be crawled.
changeDetectingColumns	Enter the names of the columns that Amazon Q will use to detect content changes. Amazon Q will re-index content when there is a change in any of these columns
allowedUsersColumns	Enter the name of the column which contains User IDs to be allowed access to content.
allowedGroupsColumn	Enter the name of the column which contains User IDs to be allowed access to content.
sourceURIColumn	Enter the name of the column which contains Source URLs to be indexed.
isSslEnabled	true to add a path to an SSL certificate file stored in an Amazon S3 bucket.
type	The type of data source. Specify JDBC as your data source type.

Configuration	Description
syncMode	<p>Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose</p> <ul style="list-style-type: none"> • <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index • <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index • <code>CHANGE_LOG</code> to incrementally crawl only new and modified content each time your data source syncs with your index.
secretArn	<p>The Amazon Resource Name (ARN) of a Secrets Manager secret that contains user name and password required to connect to your database. The secret must contain a JSON structure with the following keys:</p> <pre data-bbox="829 1247 1507 1440"> { "user name": "<i>database user name</i>", "password": "<i>password</i>" } </pre>
version	<p>The version of the template that is currently supported.</p>

How Amazon Q Business connector crawls Amazon RDS (PostgreSQL) ACLs

When you connect a database data source to Amazon Q, Amazon Q crawls user and group information from a column in the source table. You specify this column in the console or using the configuration parameter as part of the `CreateDataSource` operation.

If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

A database data source has the following limitations:

- You can only specify an allow list for a database data source. You can't specify a deny list.
- You can only specify groups. You can't specify individual users for the allow list.
- The database column should be a string containing a semicolon delimited list of groups.

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business Amazon RDS (PostgreSQL) data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a attribute mapping already available, or if you want to map additional document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q PostgreSQL connector supports the following field mappings:

Supported field mappings

- [Document](#)

Document

JDBC field name	Index field name	Description	Data type
jd_document_id	jd_document_id	Custom	String
jd_document_title	jd_document_title	Custom	String
jd_source_uri	_source_uri	Default	String

IAM role for Amazon Q Business Amazon RDS (PostgreSQL) connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the BatchPutDocument and BatchDeleteDocument operations to ingest documents.

- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- Permission to access the SSL certificate stored in your Amazon S3 bucket.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AllowsAmazonQToGetS3Objects",
    "Action": [
      "s3:GetObject"
    ],
    "Resource": [
      "arn:aws:s3:::{{input_bucket_name}}/*"
    ],
    "Effect": "Allow",
    "Condition": {
      "StringEquals": {
        "aws:ResourceAccount": "{{account_id}}"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToGetSecret",
    "Effect": "Allow",
    "Action": [
      "secretsmanager:GetSecretValue"
    ],
    "Resource": [
      "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToDecryptSecret",
    "Effect": "Allow",
    "Action": [
      "kms:Decrypt"
    ],
    "Resource": [
```

```

        "arn:aws:kms:{{region}}:{{account_id}}:key/[[key_id]]"
    ],
    "Condition": {
        "StringLike": {
            "kms:ViaService": [
                "secretsmanager.*.amazonaws.com"
            ]
        }
    }
},
{
    "Sid": "AllowsAmazonQToIngestDocuments",
    "Effect": "Allow",
    "Action": [
        "qbusiness:BatchPutDocument",
        "qbusiness:BatchDeleteDocument"
    ],
    "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
},
{
    "Sid": "AllowsAmazonQToIngestPrincipalMapping",
    "Effect": "Allow",
    "Action": [
        "qbusiness:PutGroup",
        "qbusiness:CreateUser",
        "qbusiness>DeleteGroup",
        "qbusiness:UpdateUser",
        "qbusiness:ListGroup"
    ],
    "Resource": [
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}",
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}",
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}/data-source/*"
    ]
},
{
    "Sid": "AllowsAmazonQToCreateAndDeleteNI",
    "Effect": "Allow",
    "Action": [
        "ec2:CreateNetworkInterface",

```

```

        "ec2:DeleteNetworkInterface"
    ],
    "Resource": [
        "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
        "arn:aws:ec2:{{region}}:{{account_id}}:security-group/
[[security_group]]"
    ]
},
{
    "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
    "Effect": "Allow",
    "Action": [
        "ec2:CreateNetworkInterface",
        "ec2:DeleteNetworkInterface"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
        "StringLike": {
            "aws:RequestTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
        },
        "ForAllValues:StringEquals": {
            "aws:TagKeys": [
                "AMAZON_Q"
            ]
        }
    }
},
{
    "Sid": "AllowsAmazonQToCreateTags",
    "Effect": "Allow",
    "Action": [
        "ec2:CreateTags"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
        "StringEquals": {
            "ec2:CreateAction": "CreateNetworkInterface"
        }
    }
},
{
    "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
    "Effect": "Allow",

```

```

    "Action": [
      "ec2:CreateNetworkInterfacePermission"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:ResourceTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
    "Effect": "Allow",
    "Action": [
      "ec2:DescribeNetworkInterfaces",
      "ec2:DescribeAvailabilityZones",
      "ec2:DescribeNetworkInterfaceAttribute",
      "ec2:DescribeVpcs",
      "ec2:DescribeRegions",
      "ec2:DescribeNetworkInterfacePermissions",
      "ec2:DescribeSubnets"
    ],
    "Resource": "*"
  }
]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToAssumeRoleForServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "{{source_account}}"
        }
      }
    }
  ]
}

```



```
    },
    "ArnLike": {
      "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
    }
  }
}
]
```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Known limitations for the Amazon Q Business Amazon RDS (PostgreSQL) connector

- Deleted database rows will not be tracked in when Amazon Q checks for updated content.
- The size of field names and values in a row of your database can't exceed 400KB.
- If you have a large amount of data in your database data source, and do not want Amazon Q to index all your database content after the first sync, you can choose to sync only new, modified, or deleted documents.

Connecting Amazon S3 to Amazon Q Business

Amazon Simple Storage Service (Amazon S3) is an object storage service that stores data as objects within storage buckets. You can connect an Amazon S3 instance to Amazon Q Business—using either the AWS Management Console or the [CreateDataSource](#) API—and create an Amazon Q web experience.

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).
- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [Amazon S3 connector overview](#)
- [Prerequisites for connecting Amazon Q Business to Amazon S3](#)
- [Connecting Amazon Q Business to Amazon S3 using the console](#)
- [Connecting Amazon Q Business to Amazon S3 using APIs](#)
- [Adding document metadata in Amazon S3](#)
- [How Amazon Q Business connector crawls Amazon S3 ACLs](#)
- [Amazon Q Business Amazon S3 data source connector field mappings](#)
- [IAM role for Amazon Q Business Amazon S3 connector](#)
- [Known limitations for the Amazon Q Business Amazon S3 connector](#)
- [Troubleshooting your Amazon Q Amazon S3 connector](#)

Amazon S3 connector overview

The following table gives an overview of the Amazon Q Business Amazon S3 connector and its supported features.

Category	Feature	Support
Security	Authentication type	https://docs.aws.amazon.com/amazonq/latest/qbusiness-ug/s3-iam-role.html IAM role
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Identity crawling	No
	VPC	Yes
Crawl features	Custom metadata	Yes
	Entities	Yes. The following entities are supported: <ul style="list-style-type: none"> • Document
	Field mappings	Yes. Supports both default and custom field mappings. For more information, see Field mappings .

Category	Feature	Support
	Filters	Yes. The following filters are supported: <ul style="list-style-type: none"> • Include/exclude by prefix • Include/exclude by glob patterns • Include/exclude by file types
	Sync mode	Supports full and incremental sync.
	File types	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to Amazon S3

Before you begin, make sure that you have completed the following prerequisites.

In Amazon S3, make sure you have:

- Copied the name of your Amazon S3 bucket name.

Note

Your bucket must be in the same AWS Region as your Amazon Q index, and your index must have permissions to access the bucket that contains your documents.

- If using Amazon VPC with Amazon S3 connector, make sure that you have assigned an Amazon S3 endpoint to your virtual private cloud (VPC). For more information about configuring an Amazon S3 connector with Amazon VPC, see [Using Amazon VPC with Amazon S3](#).

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your Amazon S3 authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Connecting Amazon Q Business to Amazon S3 using the console

The following procedure outlines how to connect Amazon Q Business to Amazon S3 using the AWS Management Console.

Connecting Amazon Q to Amazon S3

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **Amazon S3** page, enter the following information:
6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. **Configure VPC and security group – optional** – You can choose to use a VPC if your Amazon S3 bucket is not accessible through the public internet. If you so, you must add **Subnets** and **VPC security groups** as well.

Important

Make sure you have:

- Configured your VPC according to the steps in [Gateway endpoints for Amazon S3](#).

- Chosen a private subnet in an Amazon Q [supported availability zone](#).
- Configured your security group to allow Amazon Q to access the Amazon S3 endpoint.

For more information, see [Using Amazon VPC](#) and [Using Amazon VPC with Amazon S3](#).

If you choose to use VPC, enter the following information:

- Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
 - VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.
8. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

 **Note**

IAM roles used for applications can't be used for data sources. If you are unsure if an existing role is used for an application, choose **Create a new role** to avoid errors.

9. **Sync scope**, enter the following information:
- Enter the data source location** – The path to the Amazon S3 bucket where your data is stored. Select **Browse S3** to find and choose your bucket.
 - Maximum file size - optional** – The maximum file size value that Amazon Q will crawl. Amazon Q will only crawl files within the limit you define.
 - Advanced settings**, enter the following information:
 - **Metadata files prefix folder location - optional** – The path to the folder in which your metadata is stored. Select **Browse S3** to locate your metadata folder.
 - **Access control list configuration file location - optional** – The path to the location of a file containing a JSON structure that specifies access settings for the files stored in your S3 data source. Select **Browse S3** to locate your ACL file.

- d. **Regex patterns** – Add patterns to include or exclude documents from your index. All paths are relative to the data source location Amazon S3 bucket. You can add up to 100 patterns.

You can include and exclude documents using file names, file types, file paths, and glob patterns (patterns that can expand a wildcard pattern into a list of path names that match the given pattern).

Examples of glob patterns include:

- `/myapp/config/*` – All files inside config directory
 - `/**/* .png` – All .png files in all directories
 - `/**/*.{png,ico,md}` – All .png, .ico, or .md files in all directories
 - `/myapp/src/**/* .ts` – All .ts files inside src directory (and all its subdirectories)
 - `**/!(*.module).ts` – All .ts files but not .module.ts
10. **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.
 - **Full sync** – Sync all content regardless of the previous sync status.
 - **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.
 11. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
 12. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
 13. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:
 - a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
 - b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

Note

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

14. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

15. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

Note

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to Amazon S3 using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the `configuration` parameter to provide a JSON schema with all other configuration information specific to your data source connector.

For an example of the API request, see [CreateDataSource](#) in the Amazon Q API Reference.

Amazon S3 JSON schema

The following is the Amazon S3 JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
```

```
"properties": {
  "connectionConfiguration": {
    "type": "object",
    "properties": {
      "repositoryEndpointMetadata": {
        "type": "object",
        "properties": {
          "BucketName": {
            "type": "string"
          }
        },
        "required": [
          "BucketName"
        ]
      },
      "required": [
        "repositoryEndpointMetadata"
      ]
    },
    "required": [
      "repositoryEndpointMetadata"
    ]
  },
  "repositoryConfigurations": {
    "type": "object",
    "properties": {
      "document": {
        "type": "object",
        "properties": {
          "fieldMappings": {
            "type": "array",
            "items": [
              {
                "type": "object",
                "properties": {
                  "indexFieldName": {
                    "type": "string"
                  },
                  "indexFieldType": {
                    "type": "string",
                    "enum": [
                      "STRING"
                    ]
                  },
                  "dataSourceFieldName": {
                    "type": "string"
                  }
                }
              }
            ]
          }
        }
      }
    }
  }
}
```



```
        },
        "required": [
            "indexFieldName",
            "indexFieldType",
            "dataSourceFieldName"
        ]
    }
]
}
},
"required": [
    "fieldMappings"
]
}
},
"required": [
    "document"
]
},
"additionalProperties": {
    "type": "object",
    "properties": {
        "inclusionPatterns": {
            "type": "array"
        },
        "exclusionPatterns": {
            "type": "array"
        },
        "inclusionPrefixes": {
            "type": "array"
        },
        "exclusionPrefixes": {
            "type": "array"
        },
        "aclConfigurationFilePath": {
            "type": "string"
        },
        "metadataFilesPrefix": {
            "type": "string"
        },
        "maxFileSizeInMegaBytes": {
            "type": "string"
        }
    }
}
```

```

    },
    "syncMode": {
      "type": "string",
      "enum": [
        "FULL_CRAWL",
        "FORCED_FULL_CRAWL"
      ]
    },
    "type": {
      "type": "string",
      "pattern": "S3"
    },
    "version": {
      "type": "string",
      "anyOf": [
        {
          "pattern": "1.0.0"
        }
      ]
    }
  },
  "required": [
    "connectionConfiguration",
    "type",
    "syncMode",
    "repositoryConfigurations"
  ]
}

```

The following provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.
repositoryEndpointMetadata	The endpoint information for the data source.
BucketName	The name of your Amazon S3 bucket.

Configuration	Description
<code>repositoryConfigurations</code>	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings.
<code>additionalProperties</code>	Additional configuration options for your content in your data source
<ul style="list-style-type: none">• <code>inclusionPatterns</code>• <code>exclusionPatterns</code>• <code>inclusionPrefixes</code>• <code>exclusionPrefixes</code>	A list of regular expression patterns to include or exclude specific files in your Amazon S3 data source. Files that match the patterns are included in the index. Files that don't match the patterns are excluded from the index. If a file matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence and the file isn't included in the index.
<code>aclConfigurationFilePath</code>	The path to the file that controls access control information for your documents in an Amazon Q index.
<code>metadataFilesPrefix</code>	The location, in your Amazon S3 bucket, of your document metadata files.
<code>maxFileSizeInMegaBytes</code>	The maximum size, in megabytes, of a file that can be added to your Amazon Q index.

Configuration	Description
syncMode	Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose from the following options: <ul style="list-style-type: none">• Use <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index• Use <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index
type	The type of data source. Specify S3 as your data source type.
version	The version of the template that's supported.

Adding document metadata in Amazon S3

To customize chat results for your end users, you can add metadata to documents in an Amazon S3 bucket by using a metadata file. Metadata is additional information about a document, such as its title and the date and time it was created.

Note

For more information about how document attributes can help you customize chat results for your end users, see [???](#).

Each metadata file is associated with an indexed document. Your metadata files must be stored in the same S3 bucket as your indexed files. You can specify a location within the S3 bucket for your metadata files by using the AWS Management Console. Or, you can use the `metadataFilesPrefix` field of the Amazon S3 configuration parameter using the JSON

schema when you create an Amazon S3 data source. If you don't specify an Amazon S3 prefix, your metadata files must be stored in the same location as your indexed documents.

If you specify an Amazon S3 prefix for your metadata files, they are in a directory structure parallel to your indexed documents. Amazon Q looks only in the specified directory for your metadata. If the metadata isn't read, check that the directory location matches the location of your metadata.

The following examples show how the indexed document location maps to the metadata file location. The document's Amazon S3 key is appended to the metadata's Amazon S3 prefix and then suffixed with `.metadata.json` to form the metadata file's Amazon S3 path. The combined Amazon S3 key, the metadata's Amazon S3 prefix, and the `.metadata.json` suffix must be no more than a total of 1,024 characters. We recommend that your Amazon S3 key is less than 1,000 characters to account for additional characters when combining your key with the prefix and suffix.

```
Bucket name:
  s3://bucketName
Document path:
  documents
Metadata path:
  none
File mapping
  s3://bucketName/documents/file.txt ->
  s3://bucketName/documents/file.txt.metadata.json
```

```
Bucket name:
  s3://bucketName
Document path:
  documents/legal
Metadata path:
  metadata
File mapping
  s3://bucketName/documents/legal/file.txt ->
  s3://bucketName/metadata/documents/legal/file.txt.metadata.json
```

Your document metadata is defined in a JSON file. The file must be a UTF-8 text file without a BOM marker. The file name of the JSON file must be `<document>.<extension>.metadata.json`. In this example, `document` is the name of the document that the metadata applies to and `extension` is the file extension for the document. The document ID must be unique in `<document>.<extension>.metadata.json`.

The content of the JSON file uses the following template.

```
{
  "DocumentId": "document ID",
  "Attributes": {
    "_category": "document category",
    "_created_at": "ISO 8601 encoded string",
    "_last_updated_at": "ISO 8601 encoded string",
    "_source_uri": "document URI",
    "_version": "file version",
    "_view_count": number of times document has been viewed,
    "custom attribute key": "custom attribute value",
    additional custom attributes
  },
  "AccessControlList": [
    {
      "Name": "user name",
      "Type": "GROUP | USER",
      "Access": "ALLOW | DENY"
    }
  ],
  "Title": "document title",
  "ContentType": "For example HTML | PDF"
}
```

All of the attributes and fields are optional, so it's not necessary to include all attributes. However, you must provide a value for each attribute that you want to include; the value can't be empty. If you don't specify the `_source_uri`, the links returned by Amazon Q in the chat results point to the Amazon S3 bucket that contains the document.

Note

For information about supported document types, see [Supported document types](#).

The `_created_at` and `_last_updated_at` metadata fields are ISO 8601 encoded dates. For example, 2012-03-25T12:30:10+01:00 is the ISO 8601 date-time format for March 25, 2012, at 12:30PM (plus 10 seconds) in the Central European Time time zone.

You can add additional information to the `Attributes` field about a document that you use to filter queries or to group query responses.

You can use the `AccessControlList` field to filter the response from a query. This way, only certain users and groups have access to documents.

How Amazon Q Business connector crawls Amazon S3 ACLs

You add access control information to a document in an Amazon S3 data source using a metadata file associated with the document. You specify the file using the console or as the `aclConfigurationFilePath` parameter when you call the `CreateDataSource` or `UpdateDataSource` API and use the `configuration` parameter.

The configuration file contains a JSON structure that identifies an Amazon S3 prefix and lists the access settings for the prefix. The prefix can be a path, or it can be an individual file. If the prefix is a path, the access settings apply to all of the files in that path.

You provide three pieces of information in the file:

- The access that the entity should have. You can use `ALLOW` or `DENY`.
- The type of entity. You can use `USER` or `GROUP`.
- The name of the entity.

The JSON structure for the configuration file must be in the following format:

```
[
  {
    "keyPrefix": "s3://BUCKETNAME/prefix1/",
    "aclEntries": [
      {
        "Name": "user1",
        "Type": "USER",
        "Access": "ALLOW"
      },
      {
        "Name": "group1",
        "Type": "GROUP",
        "Access": "DENY"
      }
    ]
  },
  {
    "keyPrefix": "s3://BUCKETNAME/prefix2/",
    "aclEntries": [
```

```
[
  {
    "Name": "user2",
    "Type": "USER",
    "Access": "ALLOW"
  },
  {
    "Name": "user1",
    "Type": "USER",
    "Access": "DENY"
  },
  {
    "Name": "group1",
    "Type": "GROUP",
    "Access": "DENY"
  }
]
```

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business Amazon S3 data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q Business enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your

data source doesn't have a attribute mapping already available, or if you want to map additional document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q Amazon S3 connector supports the following entities and the associated reserved and custom attributes.

Supported entities and field mappings

- [Document](#)

Document

Amazon S3 field name	Index field name	Description	Data type
s3_document_id	s3_document_id	Default	String

IAM role for Amazon Q Business Amazon S3 connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

When you use an Amazon S3 bucket as a data source, you must provide a role that has permissions to:

- Access your Amazon S3 bucket.
- Permission to access the [BatchPutDocument](#) and [BatchDeleteDocument](#) API operations in order to ingest documents.
- Permission to access the Principal Store APIs needed to ingest access control and identity information from documents.

To allow Amazon Q to use an Amazon S3 bucket as a data source, use the following role policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToGetObjectfromS3",
      "Action": [
        "s3:GetObject"
      ],
      "Resource": [
        "arn:aws:s3:::{{input_bucket_name}}/*"
      ],
      "Effect": "Allow",
      "Condition": {
        "StringEquals": {
          "aws:ResourceAccount": "{{account_id}}"
        }
      }
    },
    {
      "Sid": "AllowsAmazonQToListS3Buckets",
      "Action": [
        "s3:ListBucket"
      ],
      "Resource": [
        "arn:aws:s3:::{{input_bucket_name}}"
      ],
      "Effect": "Allow",
      "Condition": {
        "StringEquals": {
          "aws:ResourceAccount": "{{account_id}}"
        }
      }
    }
  ],
}
```

```

{
  "Sid": "AllowsAmazonQToIngestDocuments",
  "Effect": "Allow",
  "Action": [
    "qbusiness:BatchPutDocument",
    "qbusiness:BatchDeleteDocument"
  ],
  "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
},
{
  "Sid": "AllowsAmazonQToCallPrincipalMappingAPIs",
  "Effect": "Allow",
  "Action": [
    "qbusiness:PutGroup",
    "qbusiness:CreateUser",
    "qbusiness>DeleteGroup",
    "qbusiness:UpdateUser",
    "qbusiness:ListGroups"
  ],
  "Resource": [
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}",
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}",
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}/data-source/*"
  ]
}
]
}

```

If the documents in the Amazon S3 bucket are encrypted, you must provide the following permissions to use the AWS KMS key to decrypt the documents:

```

{
  "Sid": "AllowsAmazonQToDecryptSecret",
  "Effect": "Allow",
  "Action": [
    "kms:Decrypt"
  ],
  "Resource": [
    "arn:aws:kms:{{region}}:{{account_id}}:key/[key_id]"
  ],
}

```

```

    "Condition": {
      "StringLike": {
        "kms:ViaService": [
          "secretsmanager.*.amazonaws.com"
        ]
      }
    }
  }
}

```

If you are using an Amazon VPC, you must add the following VPC access permissions to your policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToGetObjectfromS3",
      "Action": [
        "s3:GetObject"
      ],
      "Resource": [
        "arn:aws:s3:::{{input_bucket_name}}/*"
      ],
      "Effect": "Allow",
      "Condition": {
        "StringEquals": {
          "aws:ResourceAccount": "{{account_id}}"
        }
      }
    },
    {
      "Sid": "AllowsAmazonQToListS3Buckets",
      "Action": [
        "s3:ListBucket"
      ],
      "Resource": [
        "arn:aws:s3:::{{input_bucket_name}}"
      ],
      "Effect": "Allow",
      "Condition": {
        "StringEquals": {
          "aws:ResourceAccount": "{{account_id}}"
        }
      }
    }
  ]
}

```

```

    }
  },
  {
    "Sid": "AllowsAmazonQToIngestDocuments",
    "Effect": "Allow",
    "Action": [
      "qbusiness:BatchPutDocument",
      "qbusiness:BatchDeleteDocument"
    ],
    "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
    {{application_id}}/index/{{index_id}}"
  },
  {
    "Sid": "AllowsAmazonQToCallPrincipalMappingAPIs",
    "Effect": "Allow",
    "Action": [
      "qbusiness:PutGroup",
      "qbusiness:CreateUser",
      "qbusiness>DeleteGroup",
      "qbusiness:UpdateUser",
      "qbusiness:ListGroups"
    ],
    "Resource": [
      "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}",
      "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
    index/{{index_id}}",
      "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
    index/{{index_id}}/data-source/*"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateAndDeleteENI",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2>DeleteNetworkInterface"
    ],
    "Resource": [
      "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
      "arn:aws:ec2:{{region}}:{{account_id}}:security-group/[{{security_group}}]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateDeleteENI",

```

```

    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2>DeleteNetworkInterface"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:RequestTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
      },
      "ForAllValues:StringEquals": {
        "aws:TagKeys": [
          "AMAZON_Q"
        ]
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateTags",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateTags"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringEquals": {
        "ec2:CreateAction": "CreateNetworkInterface"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterfacePermission"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:ResourceTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
      }
    }
  }
},
{

```

```

    "Sid": "AllowsAmazonQToConnectToVPC",
    "Effect": "Allow",
    "Action": [
      "ec2:DescribeNetworkInterfaces",
      "ec2:DescribeAvailabilityZones",
      "ec2:DescribeNetworkInterfaceAttribute",
      "ec2:DescribeVpcs",
      "ec2:DescribeRegions",
      "ec2:DescribeNetworkInterfacePermissions",
      "ec2:DescribeSubnets"
    ],
    "Resource": "*"
  }
]
}

```

To allow Amazon Q to assume a role, use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToAssumeRoleForServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "{{source_account}}"
        },
        "ArnLike": {
          "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
        }
      }
    }
  ]
}

```

Known limitations for the Amazon Q Business Amazon S3 connector

The Amazon Q Business Amazon S3 connector has the following known limitations:

- The Amazon S3 bucket must be in the same AWS Region as your Amazon Q index, and your index must have permissions to access the bucket that contains your documents.

Troubleshooting your Amazon Q Amazon S3 connector

The following table provides information about error codes you may see for the Amazon S3 connector and suggested troubleshooting actions.

Error code	Error message	Suggested resolution
S3-5001	Profile name cannot be null or empty. Try again with a valid profile name.	Provide a valid profile name in the configuration.
S3-5002	Default AWS profile was not found. Verify the credentials file and try again.	Configure the AWS profile in the environment using "aws configure" command.
S3-5100	Bucket cannot be null or empty. Try again with a valid bucket.	Provide a valid bucket name in configuration.
S3-5101	The bucket does not exist, or it is from another region. Try again with a valid bucket.	Provide a valid bucket name that exists in the same region as the profile that is configured in the environment.
S3-5102	The ACL file is not found in the given path. Verify and try again.	Provide a valid ACL file location in configuration.
S3-5103	The ACL file reading was unsuccessful due to malformed JSON	Verify the content of ACL file. It could contain malformed JSON content.

Error code	Error message	Suggested resolution
	content. Verify and try again.	
S3-5104	Metadata file contained malformed JSON content.	Verify content of metadata files. It could contain malformed JSON content.
S3-5105	IndexFieldName cannot be null or empty.	IndexFieldName in Field Mappings should not be null or empty.
S3-5106	IndexFieldType cannot be null or empty.	IndexFieldType in Field Mappings should not be null or empty.
S3-5107	DataSourceFieldName cannot be null or empty.	DataSourceFieldName in Field Mappings should not be null or empty.
S3-5108	Only String, String List, Date and Long formats are supported for field mappings.	IndexFieldType in field mapping could contain an unsupported type. Verify field mappings and try again.
S3-5110	Unable to connect with provided Amazon S3 bucket.	Try again with valid bucket.
S3-5111	Unable to connect with provided Amazon S3 bucket due to connection timeout.	Check if the provided bucket is valid, credentials are valid, an IAM role with correct permissions has been provided, or if the VPC configuration of the data source is correct.
S3-5200	Object was not accessible. Amazon S3 returned following error:	The object might be not accessible. User may receive this error if an object is encrypted using an SSE-KMS key that the profile doesn't have the access.

Error code	Error message	Suggested resolution
S3-5201	The object content was not readable. S3 returned following error:	User may receive this error if an object is encrypted using an SSE-C key.

Connecting Amazon Q custom connector to Amazon Q Business

Use a custom data source when you have a repository that Amazon Q Business doesn't yet provide a data source connector for. When you create a custom data source, you have complete control over how the documents to index are selected. Amazon Q only provides metric information that you can use to monitor your data source sync jobs. You must create and run the crawler that determines the documents your data source indexes.

You can use a custom data source connector to:

- See the same run history metrics that Amazon Q data sources provide even when you can't use Amazon Q data sources to sync your repositories.
- Create a consistent sync monitoring experience between Amazon Q data sources and custom data sources.
- See sync metrics for a data source connector that you created using the [BatchPutDocument](#) and [BatchDeleteDocument](#) API operations.

You can create an Amazon Q custom data source connector using either the AWS Management Console or the [CreateDataSource](#).

When you create a custom data source using the CreateDataSource API operation:

- The action returns an ID to use when you synchronize the data source.
- You have to set the Configuration parameter as the following:

```
"configuration": {  
  "type": "CUSTOM",  
  "version": "1.0.0"  
}
```

- You must specify the main title of your documents using the [Document](#) object, and `_source_uri` in [DocumentAttribute](#). The main title is required so that `DocumentTitle` and `DocumentURI` are included in the [ChatSync](#) response.

When you create a custom data source using the console:

- The console returns an ID to use when you synchronize the data source.
- Give your data source a name, and optionally a description and resource tags.
- After the data source is created, a data source ID is shown. Copy this ID to use when you synchronize the data source with the index.

Topics

- [Creating an Amazon Q custom connector](#)
- [Required attributes](#)
- [Viewing metrics](#)

Creating an Amazon Q custom connector

To use a custom data source, create an application that is responsible for updating your Amazon Q index. The application depends on a crawler that you create. The crawler reads the documents in your repository and determines which documents should be sent to Amazon Q. Your application should perform the following steps:

1. Crawl your repository and make a list of the documents in your repository that are added, updated, or deleted.
2. Call the [StartDataSourceSyncJob](#) API operation to signal that a sync job is starting. You provide a data source ID to identify the data source that is synchronizing. Amazon Q returns an execution ID to identify a particular sync job.

Note

After you end a sync job, you can start a new sync job. There can be a period of time before all of the submitted documents are added to the index. To see the status of the sync job, use the [ListDataSourceSyncJobs](#) operation. If the Status returned for the sync

job is SYNCING_INDEXING, some documents are still being indexed. You can start a new sync job when the status of the previous job is FAILED or SUCCEEDED.

3. To remove documents from the index, use the [BatchDeleteDocument](#) operation. You provide the data source ID and execution ID to identify the data source that is synchronizing and the job that this update is associated with.
4. To signal the end of the sync job, use the [StopDataSourceSyncJob](#) operation. After you call the StopDataSourceSyncJob operation, the associated execution ID is no longer valid.

Note

After you call the StopDataSourceSyncJob operation, you can't use a sync job identifier in a call to the BatchPutDocument or BatchDeleteDocument operations. If you do, all of the documents submitted are returned in the FailedDocuments response message from the API.

5. To list the sync jobs for the data source and to see metrics for the sync jobs, use the [ListDataSourceSyncJobs](#) operation with the index and data source identifiers.

Required attributes

When you submit a document to Amazon Q using the BatchPutDocument API operation, you must provide the following two attributes for each document:

- `_data_source_id` – The identifier of the data source. This is returned when you create the data source with either the console or the CreateDataSource API operation.
- `_data_source_sync_job_execution_id` – The identifier of the sync run. This is returned when you start the index synchronization with the StartDataSourceSyncJob operation.

The following is the JSON required to index a document using a custom data source.

```
{
  "Documents": [
    {
      "Attributes": [
        {
          "Key": "_data_source_id",
          "Value": {
```

```

        "StringValue": "data source identifier"
      }
    },
    {
      "Key": "_data_source_sync_job_execution_id",
      "Value": {
        "StringValue": "sync job identifier"
      }
    }
  ],
  "Blob": "document content",
  "ContentType": "content type",
  "Id": "document identifier",
  "Title": "document title"
}
],
"IndexId": "index identifier",
"RoleArn": "IAM role ARN"
}

```

When you remove a document from the index using the `BatchDeleteDocument` API operation, you must specify the following two fields in the `DataSourceSyncJobMetricTarget` parameter:

- `DataSourceId` – The identifier of the data source. This is returned when you create the data source with either the console or the `CreateDataSource` API operation.
- `DataSourceSyncJobId` – The identifier of the sync run. This is returned when you start the index synchronization with the `StartDataSourceSyncJob` operation.

The following is the JSON required to delete a document from the index using the `BatchDeleteDocument` operation.

```

{
  "DataSourceSyncJobMetricTarget": {
    "DataSourceId": "data source identifier",
    "DataSourceSyncJobId": "sync job identifier"
  },
  "DocumentIdList": [
    "document identifier"
  ],
  "IndexId": "index identifier"
}

```

Viewing metrics

After a sync job is finished, you can use the `DataSourceSyncJobMetrics` API operation to get the metrics associated with the sync job. Use this API operation to monitor your custom data source syncs.

You can submit the same document multiple times, either as part of the `BatchPutDocument` operation, the `BatchDeleteDocument` operation, or if the document is submitted for both addition and deletion. Regardless of how you submit the document, it is only counted once in the metrics.

- `DocumentsAdded` – The number of documents submitted using the `BatchPutDocument` operation associated with this sync job that are added to the index for the first time. If a document is submitted for addition more than once in a sync, the document is only counted once in the metrics.
- `DocumentsDeleted` – The number of documents submitted using the `BatchDeleteDocument` operation associated with this sync job that are deleted from the index. If a document is submitted for deletion more than once in a sync, the document is only counted once in the metrics.
- `DocumentsFailed` – The number of documents associated with this sync job that failed indexing. These documents were accepted by Amazon Q for indexing but could not be indexed or deleted. If a document isn't accepted by Amazon Q, the identifier for the document is returned in the `FailedDocuments` response property of the `BatchPutDocument` and `BatchDeleteDocument` operations.
- `DocumentsModified` – The number of modified documents submitted using the `BatchPutDocument` operation associated with this sync job that were modified in the Amazon Q index.

Amazon Q also emits Amazon CloudWatch metrics while indexing documents. For more information, see [Monitoring Amazon Q with Amazon CloudWatch](#).

Amazon Q doesn't return the `DocumentsScanned` metric for custom data sources.

Connecting Web Crawler to Amazon Q Business

An Amazon Q Business Web Crawler connector crawls and indexes either public facing websites or internal company websites that use HTTPS. With Amazon Q web crawler, you can create a

generative AI web experience for your end users based on the website data you crawl using either the AWS Management Console or the [CreateDataSource](#) API.

Note

Amazon Q Web Crawler supports only HTTPS enabled sites. It doesn't support HTTP or self-signed certificate enabled websites.

Amazon Q Web Crawler uses the Selenium web crawler package and a Chromium driver. Amazon Q automatically updates the version of Selenium and the Chromium driver using continuous integration (CI).

Important

When selecting websites to index, you must adhere to the [Amazon Acceptable Use Policy](#) and all other Amazon terms. Remember that you must only use Amazon Q Web Crawler to index your own webpages, or webpages that you have authorization to index. To learn how to stop Amazon Q Web Crawler from indexing your websites, see [Configuring a robots.txt file for Amazon Q Business Web Crawler](#).

If you receive an error when crawling a website, it could be that the website is blocked from crawling. To crawl internal websites, you can set up a web proxy. The web proxy must be public facing. You can also use authentication to access and crawl websites.

Note

Amazon Q Web Crawler connector does *not* support AWS KMS encrypted Amazon S3 buckets. It supports only server-side encryption with Amazon S3 managed keys.

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).

- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [Web Crawler connector overview](#)
- [Prerequisites for connecting Amazon Q Business to Web Crawler](#)
- [Retrieving XPaths \(XML Path Language\) for Web Crawler](#)
- [Connecting Amazon Q Business to Web Crawler using the console](#)
- [Connecting Amazon Q Business to Web Crawler using APIs](#)
- [Amazon Q Business Web Crawler data source connector field mappings](#)
- [IAM role for Amazon Q Business Web Crawler connector](#)
- [Configuring a robots.txt file for Amazon Q Business Web Crawler](#)

Web Crawler connector overview

The following table gives an overview of the Amazon Q Business Web Crawler connector and its supported features.

Category	Feature	Support
Security	Authentication type	<ul style="list-style-type: none"> • Basic • NTLM/Kerberos • Form • SAML <div style="border: 1px solid #add8e6; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p>Note</p> <p>You don't need authentication to crawl public websites you have permission to crawl.</p> </div>
	Authentication credentials	<p>Basic authentication</p> <ul style="list-style-type: none"> • Website username

Category	Feature	Support
		<ul style="list-style-type: none"> Website password <p>NTLM/Kerberos authentication</p> <ul style="list-style-type: none"> NTLM/Kerberos username NTLM/Kerberos password <p>Form authentication</p> <ul style="list-style-type: none"> Login page URL Website username Website password Username field Xpath Password field Xpath Password button Xpath (Optional) Username button Xpath <p>SAML authentication</p> <ul style="list-style-type: none"> Login page URL Website username Website password Username field Xpath Password field Xpath Password button Xpath (Optional) Username button Xpath
	<u>Access Control List (ACL) crawling</u>	No
	<u>Identity crawling</u>	No
Crawl features	Custom metadata	Yes

Category	Feature	Support
	Entities	Yes. The following entities are supported: <ul style="list-style-type: none"> • Web page • Attachment
	Field mappings	Yes. For more information, see Field mappings .
	Filters	Yes. The following filters are supported : <ul style="list-style-type: none"> • Filter comments in files • Sync specific domains and subdomains • Include files linked on web pages • Regex patterns to crawl and index specific URLs • Regex patterns to crawl and index specific files • Include web pages by crawl depth • Specify maximum file size and links per page for Amazon Q to crawl
	Sync mode	Supports full and new, modified, or deleted content sync
	File types	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to Web Crawler

Before you begin, make sure that you have completed the following prerequisites.

For Amazon Q Web Crawler, make sure you have:

- Copied the seed or sitemap URLs of the websites that you want to index and stored them in a text file or an Amazon S3 bucket. Each URL must be included on a separate line.
- **For XML sitemaps:** Copied the sitemap XML and saved it in an XML file in an Amazon S3 bucket. You can also combine multiple sitemap XML files into a .zip file.
- **For websites that require basic, NTLM, or Kerberos authentication:**
 - Noted your website authentication credentials, which include a username and password.

 **Note**

Amazon Q Web Crawler supports the NTLM authentication protocol that includes password hashing, and Kerberos authentication protocol that includes password encryption.

- **For websites that require SAML or login form authentication:**
 - Noted your website authentication credentials, which include a username and password.
 - Copied the XPath(s) (XML Path Language) of the username field (and the username button if using SAML), password field and button, and copied the login page URL. You can find the XPath(s) of elements using your web browser's developer tools. XPath(s) follow this format: `//tagname[@Attribute='Value']`.

 **Note**

Amazon Q Web Crawler uses a headless Chrome browser and the information from the form to authenticate and authorize access with an OAuth 2.0 protected URL.

- **Optional:** Copied the host name and the port number of the web proxy server if you want to use a web proxy to connect to internal websites that you want to crawl. The web proxy must be public facing. Amazon Q supports connecting to web proxy servers backed by basic authentication, or you can connect with no authentication.
- **Optional:** Copied the virtual private cloud (VPC) subnet ID if you want to use a VPC to connect to internal websites you want to crawl. For more information, see [Using Amazon VPC](#).

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the API, noted the ARN of the IAM role.

- **For websites that require authentication credentials to crawl:** Stored your Web Crawler authentication credentials in an AWS Secrets Manager secret and, if using the API, noted the ARN of the secret.

 **Note**

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

Retrieving XPath (XML Path Language) for Web Crawler

If the website you are crawling with Amazon Q Business Web Crawler uses Form or SAML authentication, you need to provide Amazon Q with the absolute XPath for the username and password fields on your web page. Optionally, you may also need to provide the absolute XPath to the username and password buttons.

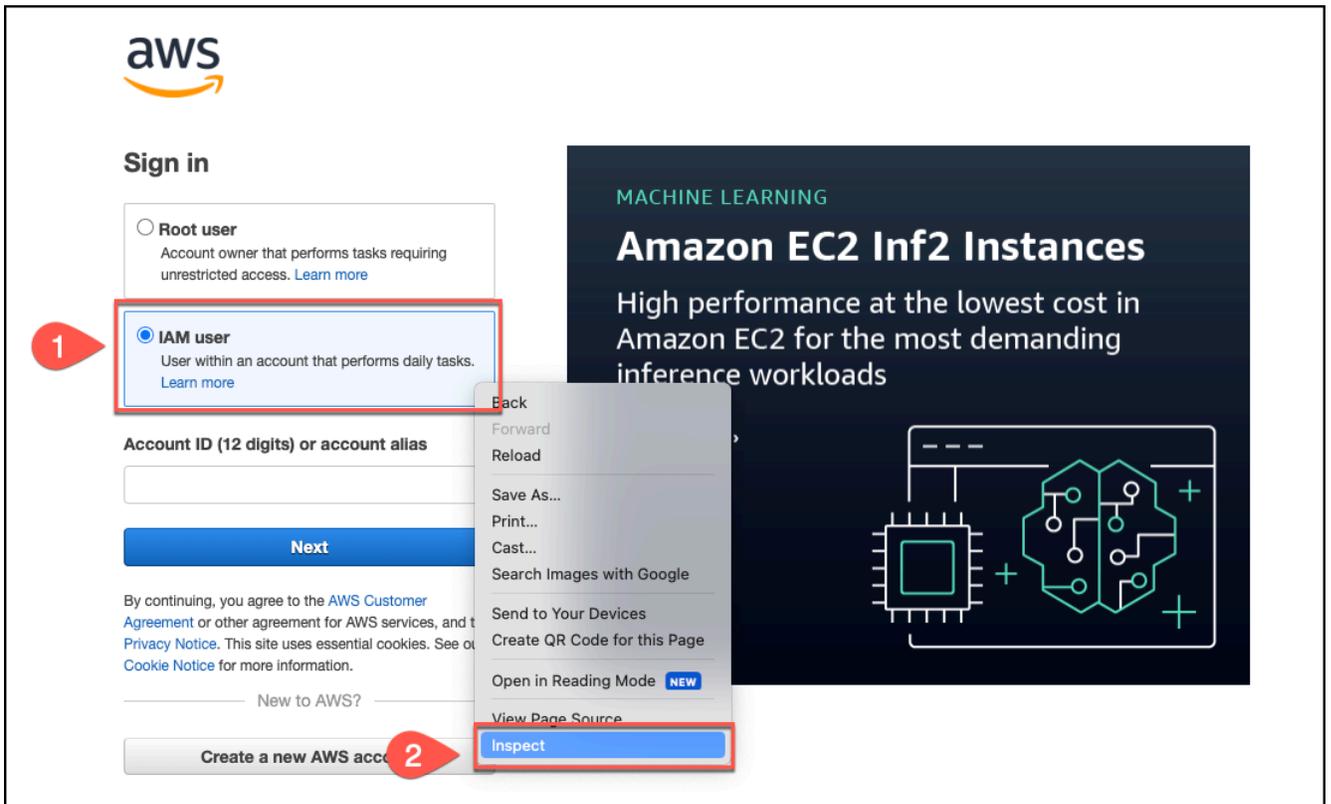
XPaths are expressions used to uniquely identify and locate the content of any XML like language document (including HTML). Amazon Q uses the XPath you provide to confirm access to the website you want to crawl. XPath usually follow the following format: `// tagname[@Attribute='Value']`.

The following tabs provide a procedure for retrieving XPath required for your Amazon Q Web Crawler connector using different web browsers.

Chrome

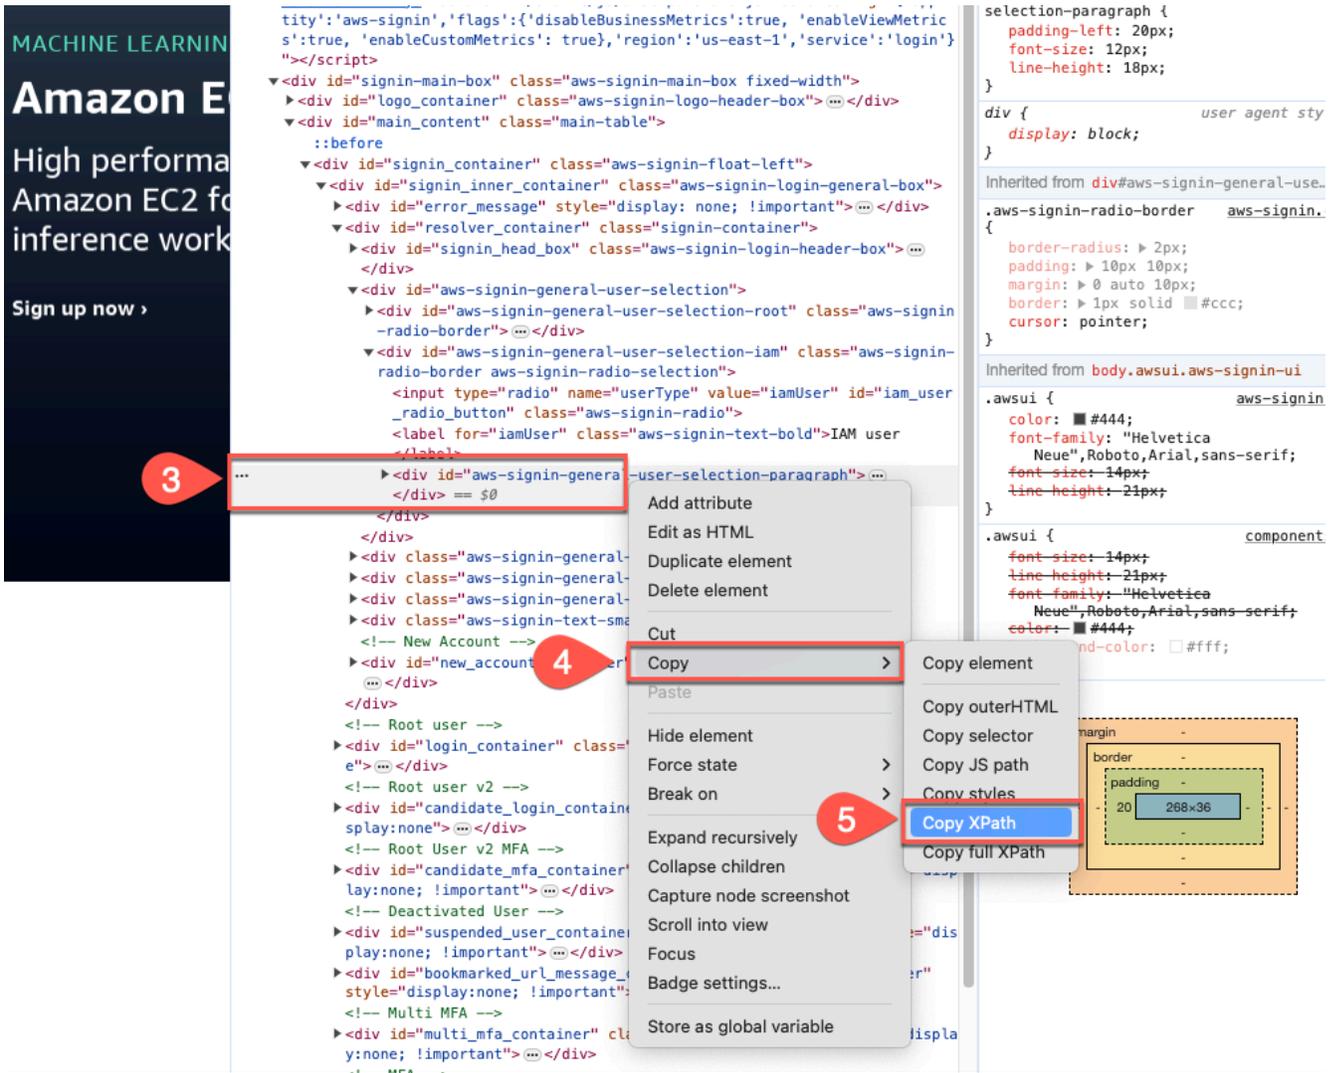
To retrieve XPath for an Amazon Q Web Crawler

1. Make sure you're on the web page you want to crawl. Then, either select or click on the web page element you want to retrieve the XPath for. This could be the username or password fields, or the username and password buttons.
2. Then, open the context (right-click) menu and then choose the **Inspect** option.



In the **Developer Tools** window that opens, the details for the element you've chosen will be highlighted.

3. Right click on the highlighted element to open the context (right-click) menu.
4. Choose **Copy**.
5. Then, choose **Copy XPath**.



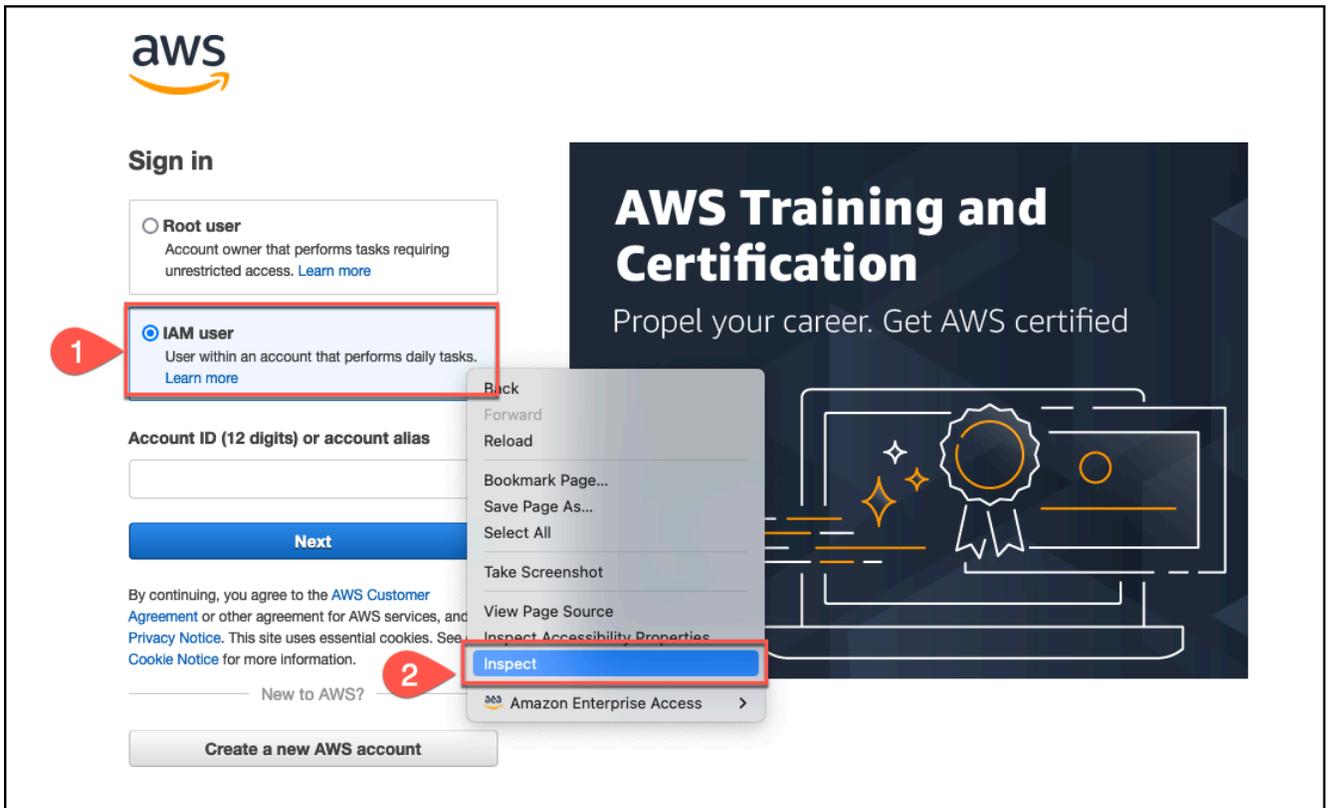
6. Then, open a text editor of your choice and paste the XPath you copied. The format of the XPath will look like this: `//tagname[@Attribute='Value']`.

Input the relevant XPaths you've copied in the **Authentication** section when you configure Amazon Q Web Crawler connector.

Firefox

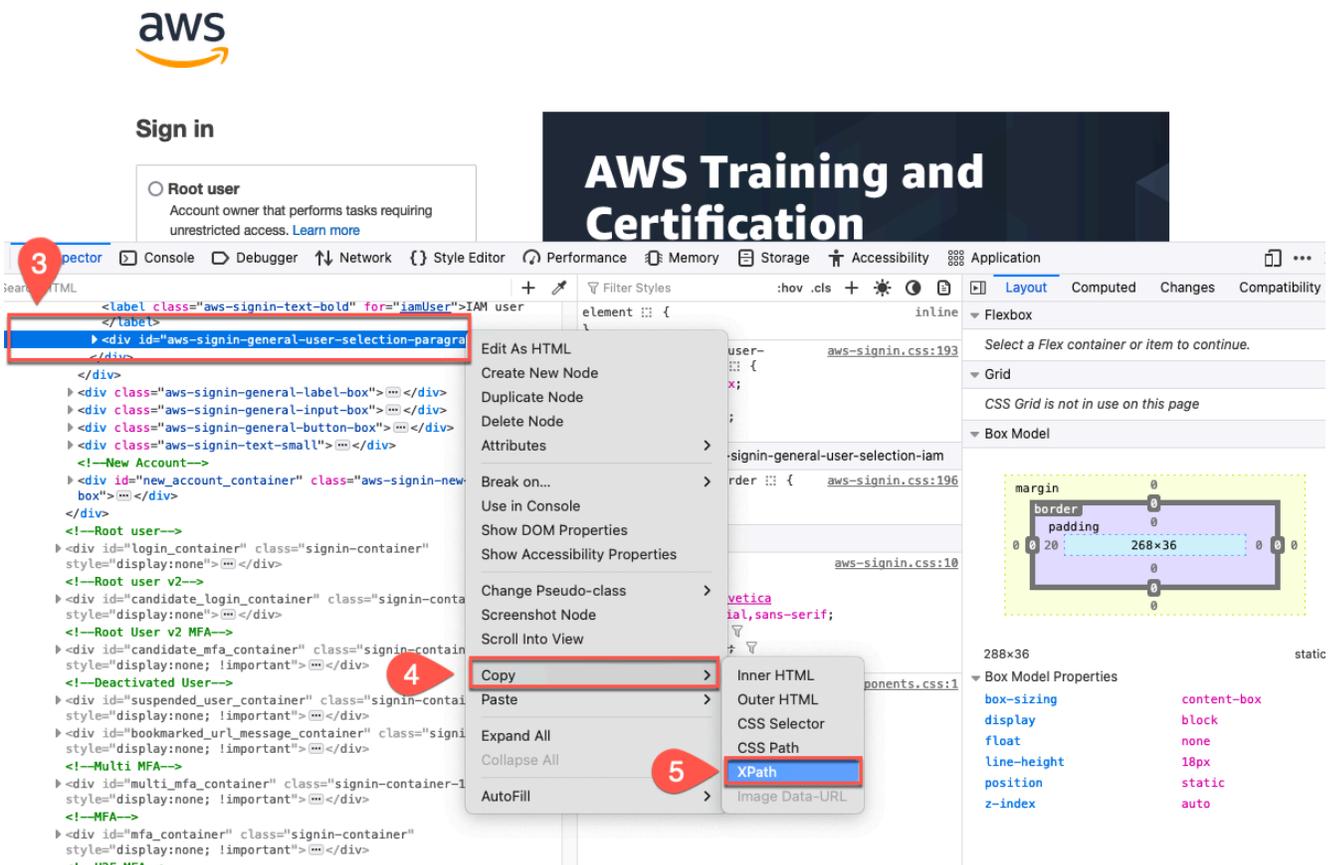
To retrieve XPaths for an Amazon Q Web Crawler

1. Make sure you're on the web page you want to crawl. Then, either select or click on the web page element you want to retrieve the XPath for. This could be the username or password fields, or the username and password buttons.
2. Then, open the context (right-click) menu and then choose the **Inspect** option.



In the **Developer Tools** window that opens, the details for the element you've chosen will be highlighted.

3. Right click on the highlighted element to open the context (right-click) menu.
4. Choose **Copy**.
5. Then, choose **Copy XPath**.



- Then, open a text editor of your choice and paste the XPath you copied. The format of the XPath will look like this: `//tagname[@Attribute='Value']`.

Input the relevant XPaths you've copied in the **Authentication** section when you configure Amazon Q Web Crawler connector.

Connecting Amazon Q Business to Web Crawler using the console

On the **Web Crawler** page, enter the following information:

- Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

- In **Source** choose from the following options:

- Source URLs** – Add up to 10 seed/starting point URLs of the websites you want to crawl. You can also include website subdomains.
- Source sitemaps** – Add up to 3 sitemap URLs of the websites you want to crawl.

- **Source URLs file** – Add up to 100 seed/starting point URLs listed in a text file in Amazon S3. Each URL should be on a separate line in the text file.
- **Source sitemaps file** – Add up to 3 sitemap XML files stored in Amazon S3. You can also zip the XML files.

 **Note**

If you choose to use a text file that includes a list of up to 100 seed URLs or to use a sitemap XML file, you specify the path to the Amazon S3 bucket where your file is stored.

You can also combine multiple sitemap XML files into a .zip file. Otherwise, you can manually enter up to 10 seed or starting point URLs, and up to three sitemap URLs.

 **Note**

If you want to crawl a sitemap, check that the base or root URL is the same as the URLs listed on your sitemap page. For example, if your sitemap URL is *https://example.com/sitemap-page.html*, the URLs listed on this sitemap page should also use the base URL "https://example.com/".

 **Note**

If you want to later edit your data source to change your seed URLs with authentication to sitemaps, you must create a new data source.

Amazon Q configures the data source using the seed URLs endpoint information in the Secrets Manager secret for authentication. Therefore, Amazon Q can't reconfigure the data source when changing to sitemaps.

3. In **Authentication**, choose the type of authentication you want to use and enter the following information in your AWS Secrets Manager secret:

- **No authentication** – Choose to crawl a public website without any authentication.
- **Basic authentication** – Enter a name for the secret, plus the username and password

- **NTLM/Kerberos authentication** – Enter a name for the secret, plus the username and password. NTLM authentication protocol includes password hashing, and Kerberos authentication protocol includes password encryption
 - **Form authentication** – Enter a name for the secret, and the username and password. Use XPath for the username field. Use XPaths for the password field and button, and login page URL. You can find the XPaths (XML Path Language) of elements using your web browser's developer tools. XPaths usually follow this format: `//tagname[@Attribute='Value']`
 - **SAML authentication** – Enter a name for the secret, plus the username and password. Use XPath for the username field and for the username button. Use XPaths for the password field and button, and login page URL. You can find the XPaths (XML Path Language) of elements using your web browser's developer tools. XPaths usually follow this format: `//tagname[@Attribute='Value']`
4. **Web proxy – optional** – Enter the host name and the port number of the proxy server that you want to use to connect to internal websites. For example, the host name of `https://a.example.com/page1.html` is "a.example.com" and the port number is 443, the standard port for HTTPS. If web proxy credentials are required to connect to a website host, you can create an AWS Secrets Manager secret that stores the credentials.
 5. **Configure VPC and security group – optional** – Choose whether you want to use a VPC. If you do, enter the following information:
 - a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
 - b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

6. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

7. In **Sync scope**, enter the following information:

- a. **Sync domain range** – Choose whether to sync website domains with subdomains only, or also crawl other domains that the webpages link to (**Sync everything**). By default, Amazon Q only syncs the domains of the websites that you want to crawl.
- b. In **Additional configuration – optional** – Configure the following settings:
 - **Scope settings**, choose from the following:
 - **Crawl depth** – The depth, or number, of levels from the seed level to crawl. For example, the seed URL page is depth 1 and any hyperlinks on this page that are also crawled are depth 2.
 - **Maximum file size** – The maximum size in MB of a webpage or attachment to crawl.
 - **Maximum links per page** – The maximum number of URLs on a single webpage to crawl.
 - **Maximum throttling** – The maximum number of URLs crawled per website host per minute.
 - **Include files that web pages link to** – Choose to crawl files that the webpages link to.
 - **Crawl URL patterns** – Add regular expression patterns to include or exclude crawling specific URLs, and indexing any hyperlinks on these URL webpages.
 - **URL pattern to index** – Add regular expression patterns to include or exclude crawling specific URLs, and indexing any hyperlinks on these URL webpages.
8. In **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.
 - **Full sync** – Sync all content regardless of the previous sync status.
 - **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.
9. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
10. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
11. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:

- a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
- b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

 **Note**

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

12. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

13. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

 **Note**

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to Web Crawler using APIs

To connect Amazon Q Business to Web Crawler using the Amazon Q API, call `CreateDataSource`. Use this API to:

- provide a name and tags for your data source
- an Amazon Resource Name (ARN) of an IAM role with permission to access the data source and required resources

- a sync schedule for Amazon Q to check the documents in your data source
- a Amazon VPC configuration

For more information on available parameters, see [CreateDataSource](#) in the [Amazon Q API reference](#).

Provide the seed or starting point URLs, or the sitemap URLs, as part of the connection configuration or repository endpoint details. Also specify the website authentication credentials and authentication type if your websites require authentication, and other necessary configurations.

Web Crawler JSON schema

The following is the Web Crawler JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
      "type": "object",
      "properties": {
        "repositoryEndpointMetadata": {
          "type": "object",
          "properties": {
            "siteMapUrls": {
              "type": "array",
              "items": {
                "type": "string",
                "pattern": "https://.*"
              }
            },
            "s3SeedUrl": {
              "type": ["string", "null"],
              "pattern": "s3:.*"
            },
            "s3SiteMapUrl": {
              "type": ["string", "null"],
              "pattern": "s3:.*"
            }
          }
        },
        "seedUrlConnections": {
          "type": "array",
```

```
    "items": [
      {
        "type": "object",
        "properties": {
          "seedUrl": {
            "type": "string",
            "pattern": "https://.*"
          }
        },
        "required": [
          "seedUrl"
        ]
      }
    ],
    "authentication": {
      "type": "string",
      "enum": [
        "NoAuthentication",
        "BasicAuth",
        "NTLM_Kerberos",
        "Form",
        "SAML"
      ]
    }
  },
  "required": [
    "repositoryEndpointMetadata"
  ],
  "repositoryConfigurations": {
    "type": "object",
    "properties": {
      "webPage": {
        "type": "object",
        "properties": {
          "fieldMappings": {
            "type": "array",
            "items": [
              {
                "type": "object",
                "properties": {
```

```

        "indexFieldName": {
          "type": "string"
        },
        "indexFieldType": {
          "type": "string",
          "enum": [
            "STRING",
            "DATE",
            "LONG"
          ]
        },
        "dataSourceFieldName": {
          "type": "string"
        },
        "dateFieldFormat": {
          "type": "string",
          "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
        }
      },
      "required": [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
      ]
    }
  ]
}
},
"required": [
  "fieldMappings"
]
},
"attachment": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": [
        {
          "type": "object",
          "properties": {
            "indexFieldName": {
              "type": "string"
            }
          }
        }
      ]
    }
  }
}

```

```

        "indexFieldType": {
          "type": "string",
          "enum": [
            "STRING",
            "DATE",
            "LONG"
          ]
        },
        "dataSourceFieldName": {
          "type": "string"
        },
        "dateFieldFormat": {
          "type": "string",
          "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
        }
      ],
      "required": [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
      ]
    }
  ]
},
"required": [
  "fieldMappings"
]
}
},
"syncMode": {
  "type": "string",
  "enum": [
    "FORCED_FULL_CRAWL",
    "FULL_CRAWL"
  ]
},
"additionalProperties": {
  "type": "object",
  "properties": {
    "rateLimit": {
      "type": "string",
      "default": "300"
    }
  }
}

```



```
    },
    "maxFileSize": {
      "type": "string",
      "default": "50"
    },
    },
    "crawlDepth": {
      "type": "string",
      "default": "2"
    },
    },
    "maxLinksPerUrl": {
      "type": "string",
      "default": "100"
    },
    },
    "crawlSubDomain": {
      "type": "boolean",
      "default": false
    },
    },
    "crawlAllDomain": {
      "type": "boolean",
      "default": false
    },
    },
    "honorRobots": {
      "type": "boolean",
      "default": false
    },
    },
    "crawlAttachments": {
      "type": "boolean",
      "default": false
    },
    },
    "inclusionURLCrawlPatterns": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    },
    "exclusionURLCrawlPatterns": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    },
    "inclusionURLIndexPatterns": {
      "type": "array",
      "items": {
```

```
    "type": "string"
  }
},
"exclusionURLIndexPatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"inclusionFileIndexPatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"exclusionFileIndexPatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"proxy": {
  "type": "object",
  "properties": {
    "host": {
      "type": "string"
    },
    "port": {
      "type": "string"
    },
    "secretArn": {
      "type": "string",
      "minLength": 20,
      "maxLength": 2048
    }
  }
}
},
"required": [
  "rateLimit",
  "maxFileSize",
  "crawlDepth",
  "crawlSubDomain",
  "crawlAllDomain",
```

```

        "maxLinksPerUrl",
        "honorRobots"
    ]
},
"type": {
    "type": "string",
    "pattern": "WEBCRAWLERV2"
},
"secretArn": {
    "type": "string",
    "minLength": 20,
    "maxLength": 2048
}
},
"version": {
    "type": "string",
    "anyOf": [
        {
            "pattern": "1.0.0"
        }
    ]
},
"required": [
    "connectionConfiguration",
    "repositoryConfigurations",
    "syncMode",
    "type",
    "additionalProperties"
]
}

```

The following provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.
repositoryEndpointMetadata	The endpoint information for the data source.

Configuration	Description
siteMapUrls	The list of sitemap URLs for the websites that you want to crawl. You can list up to three sitemap URLs.
s3SeedUrl	The S3 path to the text file that stores the list of seed or starting point URLs. For example, <i>s3://bucket-name/directory/</i> . Each URL in the text file must be formatted on a separate line. You can list up to 100 seed URLs in a file.
s3SiteMapUrl	The S3 path to the sitemap XML files. For example, <i>s3://bucket-name/directory/</i> . You can list up to three sitemap XML files. You can club together multiple sitemap files into a .zip file and store the .zip file in your Amazon S3 bucket.
seedUrlConnections	The list of seed or starting point URLs for the websites that you want to crawl. You can list up to 100 seed URLs.
seedUrl	The seed or starting point URL.
authentication	The authentication type if your websites require the same authentication, otherwise specify <code>NoAuthentication</code> .
repositoryConfigurations	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings.

Configuration	Description
<ul style="list-style-type: none"> • <code>webPage</code> • <code>attachment</code> 	<p>A list of objects that map the attributes or field names of your webpages and webpage files to Amazon Q index field names. For example, the HTML webpage title tag can be mapped to the <code>_document_title</code> index field.</p>
<code>syncMode</code>	<p>Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose between the following options:</p> <ul style="list-style-type: none"> • Use <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index. • Use <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index.
<code>additionalProperties</code>	<p>Additional configuration options for your content in your data source.</p>
<code>rateLimit</code>	<p>The maximum number of URLs crawled per website host per minute.</p>
<code>maxFileSize</code>	<p>The maximum size (in MB) of a webpage or attachment to crawl.</p>
<code>crawlDepth</code>	<p>The number of levels from the seed URL to crawl. For example, the seed URL page is depth 1 and any hyperlinks on this page that are also crawled are depth 2.</p>

Configuration	Description
<code>maxLinksPerUrl</code>	<p>The maximum number of URLs on a webpage to include when crawling a website. This number is per webpage. As a website's webpages are crawled, any URLs that the webpages link to also are crawled. URLs on a webpage are crawled in order of appearance.</p>
<code>crawlSubDomain</code>	<p><code>true</code> to crawl the website domains with subdomains only. For example, if the seed URL is "abc.example.com", then "a.abc.example.com" and "b.abc.example.com" are also crawled. If you don't set <code>crawlSubDomain</code> or <code>crawlAllDomain</code> to <code>true</code>, then Amazon Q only crawls the domains of the websites that you want to crawl.</p>
<code>crawlAllDomain</code>	<p><code>true</code> to crawl the website domains with subdomains and other domains the web pages link to. If you don't set <code>crawlSubDomain</code> or <code>crawlAllDomain</code> to <code>true</code>, then Amazon Q only crawls the domains of the websites that you want to crawl.</p>
<code>honorRobots</code>	<p><code>true</code> to respect the robots.txt directives of the websites that you want to crawl. These directives control how Amazon Q Web Crawler crawls the websites, and whether Amazon Q can crawl only specific content or not crawl any content.</p> <div data-bbox="829 1598 1507 1864" style="border: 1px solid #f08080; border-radius: 10px; padding: 10px;"><p> Important</p><p>The <code>honorRobots</code> feature is currently only available if you use the API.</p></div>

Configuration	Description
<code>crawlAttachments</code>	<p><code>true</code> to crawl files that the webpages link to.</p>
<ul style="list-style-type: none"> • <code>inclusionURLCrawlPatterns</code> • <code>inclusionURLIndexPatterns</code> 	<p>A list of regular expression patterns to <i>include</i> crawling certain URLs and indexing any hyperlinks on these URL webpages. URLs that match the patterns are included in the index. URLs that don't match the patterns are excluded from the index. If a URL matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence, and the URL and website's webpages aren't included in the index.</p>
<ul style="list-style-type: none"> • <code>exclusionURLCrawlPatterns</code> • <code>exclusionURLIndexPatterns</code> 	<p>A list of regular expression patterns to <i>exclude</i> crawling certain URLs and indexing any hyperlinks on these URL webpages. URLs that match the patterns are excluded from the index. URLs that don't match the patterns are included in the index. If a URL matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence, and the URL/website's webpages aren't included in the index.</p>
<code>inclusionFileIndexPatterns</code>	<p>A list of regular expression patterns to <i>include</i> certain web page files. Files that match the patterns are included in the index. Files that don't match the patterns are excluded from the index. If a file matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence, and the file isn't included in the index.</p>

Configuration	Description
<code>exclusionFileIndexPatterns</code>	A list of regular expression patterns to <i>exclude</i> certain webpage files. Files that match the patterns are excluded from the index. Files that don't match the patterns are included in the index. If a file matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence, and the file isn't included in the index.
<code>proxy</code>	Configuration information required to connect to your internal websites through a web proxy.
<code>host</code>	The host name of the proxy server that you want to use to connect to internal websites. For example, the host name of <code>https://a.example.com/page1.html</code> is "a.example.com".
<code>port</code>	The port number of the proxy server that you want to use to connect to internal websites. For example, 443 is the standard port for HTTPS.
<code>secretArn (proxy)</code>	If web proxy credentials are required to connect to a website host, you can create an AWS Secrets Manager secret that stores the credentials. Provide the Amazon Resource Name (ARN) of the secret.
<code>type</code>	The type of data source. Specify <code>WEBCRAWLERV2</code> as your data source type.

Configuration	Description
secretARN	<p>The Amazon Resource Name (ARN) of an AWS Secrets Manager secret that's used if your websites require authentication to access the websites. You store the authentication credentials for the website in the secret that contains JSON key-value pairs.</p> <p>If you use basic, or NTLM/Kerberos, enter the username and password. The JSON keys in the secret must be <code>username</code> and <code>password</code>. NTLM authentication protocol includes password hashing, and Kerberos authentication protocol includes password encryption.</p> <p>If you use SAML or form authentication, enter the username and password, XPath for the username field (and username button if using SAML), XPaths for the password field and button, and the login page URL. The JSON keys in the secret must be <code>username</code>, <code>password</code>, <code>usernameFieldXPath</code> , <code>usernameButtonXPath</code> , <code>passwordFieldXPath</code> , <code>passwordButtonXPath</code> , and <code>loginPageUrl</code> . You can find the XPaths (XML Path Language) of elements using your web browser's developer tools. XPaths usually follow this format: <code>//tagname[@Attribute='Value']</code> .</p> <p>Amazon Q also checks if the endpoint information (seed URLs) included in the secret is the same the endpoint information specified in your data source endpoint configuration details.</p>

Configuration	Description
version	The version of this template that's currently supported.

Amazon Q Business Web Crawler data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q Business enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a attribute mapping already available, or if you want to map additional document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q Web Crawler connector supports the following entities and the associated reserved and custom attributes.

Supported entities and field mappings

- [Web Pages](#)
- [Attachments](#)

Web Pages

Web Crawler field name	Index field name	Description	Data type
category	_category	Default	String
sourceUrl	_source_uri	Default	String
title	wc_title	Custom	String
htmlSize	wc_html_size	Custom	Long (numeric)

Attachments

Web Crawler field name	Index field name	Description	Data type
category	_category	Default	String
sourceUrl	_source_uri	Default	String
fileName	wc_file_name	Custom	String
fileType	wc_file_type	Custom	String
fileSize	wc_file_size	Custom	Long (numeric)

IAM role for Amazon Q Business Web Crawler connector

To connect Web Crawler to Amazon Q Business, you must give Amazon Q an IAM role that has the following permissions.

If you're crawling a public website with no authentication:

- Permission to access the BatchPutDocument and BatchDeleteDocument operations to ingest documents.
- Permission to access the [User Store](#) operations to ingest access control information from documents.

```
{
    "Sid": "AllowsAmazonQToIngestDocuments",
    "Effect": "Allow",
    "Action": [
        "qbusiness:BatchPutDocument",
        "qbusiness:BatchDeleteDocument"
    ],
    "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
},
{
    "Sid": "AllowsAmazonQToIngestPrincipalMapping",
    "Effect": "Allow",
    "Action": [
        "qbusiness:PutGroup",
        "qbusiness:CreateUser",
        "qbusiness>DeleteGroup",
        "qbusiness:UpdateUser",
        "qbusiness:ListGroups"
    ],
    "Resource": [
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}",
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}",
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}/data-source/*"
    ]
}
```

If you're crawling a website which uses authentication:

- Permission to access the AWS Secrets Manager secret that contains the credentials to connect to websites or a web proxy server backed by basic authentication.

```
{
  "Sid": "AllowsAmazonQToGetSecret",
  "Effect": "Allow",
  "Action": [
    "secretsmanager:GetSecretValue"
  ],
  "Resource": [
    "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
  ]
}
```

If your Secrets Manager secret is decrypted, add permissions for a AWS KMS key to decrypt the user name and password secret stored by Secrets Manager:

```
{
  "Sid": "AllowsAmazonQToDecryptSecret",
  "Effect": "Allow",
  "Action": [
    "kms:Decrypt"
  ],
  "Resource": [
    "arn:aws:kms:{{region}}:{{account_id}}:key/[key_id]"
  ],
  "Condition": {
    "StringLike": {
      "kms:ViaService": [
        "secretsmanager.*.amazonaws.com"
      ]
    }
  }
}
```

If your Amazon Q data source connector needs access to an object stored in an Amazon S3 bucket—like seed URLs or sitemaps— you must add the following permissions to your IAM role:

Note

Check that the file path to the object in your Amazon S3 bucket is of the following format:
s3://BucketName/FolderName/FileName.extension.

```
{
  "Sid": "AllowsAmazonQToGetS3Objects",
  "Action": [
    "s3:GetObject"
  ],
  "Resource": [
    "arn:aws:s3:::{{input_bucket_name}}/*"
  ],
  "Effect": "Allow",
  "Condition": {
    "StringEquals": {
      "aws:ResourceAccount": "{{account_id}}"
    }
  }
}
```

If you are using an Amazon VPC, you need to add the following VPC access permissions to your policy:

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AllowsAmazonQToCreateAndDeleteNI",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2:DeleteNetworkInterface"
    ],
    "Resource": [
      "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
      "arn:aws:ec2:{{region}}:{{account_id}}:security-group/[{{security_group}}]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2:DeleteNetworkInterface"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
```

```

    "aws:RequestTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
  },
  "ForAllValues:StringEquals": {
    "aws:TagKeys": [
      "AMAZON_Q"
    ]
  }
},
{
  "Sid": "AllowsAmazonQToCreateTags",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateTags"
  ],
  "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
  "Condition": {
    "StringEquals": {
      "ec2:CreateAction": "CreateNetworkInterface"
    }
  }
},
{
  "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateNetworkInterfacePermission"
  ],
  "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
  "Condition": {
    "StringLike": {
      "aws:ResourceTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
    }
  }
},
{
  "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
  "Effect": "Allow",
  "Action": [
    "ec2:DescribeNetworkInterfaces",
    "ec2:DescribeAvailabilityZones",
    "ec2:DescribeNetworkInterfaceAttribute",
    "ec2:DescribeVpcs",
    "ec2:DescribeRegions",

```

```

    "ec2:DescribeNetworkInterfacePermissions",
    "ec2:DescribeSubnets"
  ],
  "Resource": "*"
}
]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "{{source_account}}"
        },
        "ArnEquals": {
          "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
        }
      }
    }
  ]
}

```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Configuring a robots.txt file for Amazon Q Business Web Crawler

Amazon Q Business Web Crawler respects standard robots.txt directives like Allow and Disallow. You can modify the robot.txt file of your website to control how Amazon Q Web Crawler crawls your website.

Topics

- [Configuring how Amazon Q Web Crawler accesses your website](#)
- [Stopping Amazon Q Web Crawler from crawling your website](#)

Configuring how Amazon Q Web Crawler accesses your website

You can control how the Amazon Q Web Crawler indexes your website using Allow and Disallow directives. You can also control which web pages are indexed and which web pages are not crawled.

To allow Amazon Q Web Crawler to crawl all web pages except disallowed web pages, use the following directive:

```
User-agent: amazon-QBusiness    # Amazon Q Web Crawler
Disallow: /credential-pages/    # disallow access to specific pages
```

To allow Amazon Q Web Crawler to crawl only specific web pages, use the following directive:

```
User-agent: amazon-QBusiness    # Amazon Q Web Crawler
Allow: /pages/                  # allow access to specific pages
```

To allow Amazon Q Web Crawler to crawl all website content and disallow crawling for any other robots, use the following directive:

```
User-agent: amazon-QBusiness    # Amazon Q Web Crawler
Allow: /                         # allow access to all pages
User-agent: *                    # any (other) robot
Disallow: /                       # disallow access to any pages
```

Stopping Amazon Q Web Crawler from crawling your website

You can stop Amazon Q Web Crawler from indexing your website using the Disallow directive. You can also control which web pages are crawled and which aren't.

To stop Amazon Q Web Crawler from crawling the website, use the following directive:

```
User-agent: amazon-QBusiness    # Amazon Q Web Crawler
Disallow: /                       # disallow access to any pages
```

Amazon Q Web Crawler also supports the robots noindex and nofollow directives in meta tags in HTML pages. These directives stop the web crawler from indexing a web page and stops

following any links on the web page. You put the meta tags in the section of the document to specify the rules of robots rules.

For example, the below web page includes the directives robots noindex and nofollow:

```
<html>
<head>
  <meta name="robots" content="noindex, nofollow"/>
  ...
</head>
<body>...</body>
</html>
```

If you have any questions or concerns about Amazon Q Web Crawler, you can reach out to the [AWS support team](#).

Connecting Amazon WorkDocs to Amazon Q Business

Amazon WorkDocs is a secure content collaboration service for creating, editing, storing, and sharing content. Amazon Q Business can connect to your Amazon WorkDocs instance.

You can connect your Amazon WorkDocs instance to Amazon Q—using either the AWS Management Console, CLI, or the [CreateDataSource](#) API—and create an Amazon Q web experience.

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).
- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [Amazon WorkDocs connector overview](#)
- [Prerequisites for connecting Amazon Q Business to Amazon WorkDocs](#)
- [Connecting Amazon Q Business to Amazon WorkDocs using the console](#)
- [Connecting Amazon Q Business to Amazon WorkDocs using APIs](#)

- [How Amazon Q Business connector crawls Amazon WorkDocs ACLs](#)
- [Amazon Q Business Amazon WorkDocs data source connector field mappings](#)
- [IAM role for Amazon Q Business Amazon WorkDocs connector](#)
- [Troubleshooting your Amazon Q Business Amazon WorkDocs connector](#)

Amazon WorkDocs connector overview

The following table gives an overview of the Amazon Q Business Amazon WorkDocs connector and its supported features.

Category	Feature	Support
Security	Authentication type	https://docs.aws.amazon.com/amazonq/latest/qbusiness-ug/workdocs-connector.html#data-source-secrets-vpc-iam IAM role
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Identity crawling	No
	VPC	Yes
Crawl features	Custom metadata	Yes
	Field mappings	Yes. Supports both default and custom field mappings. For more information, see Field mappings .
	Filters	Yes. The following filters are supported: <ul style="list-style-type: none"> • Include/exclude by file name • Include/exclude by file type • Include/exclude by file path
	Sync mode	Supports full and incremental sync.
	File types	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to Amazon WorkDocs

Before you begin, make sure that you have completed the following prerequisites.

In Amazon WorkDocs, make sure you have:

- Noted the Amazon WorkDocs directory ID (organization ID) for your Amazon WorkDocs repository.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your Amazon WorkDocs authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Connecting Amazon Q Business to Amazon WorkDocs using the console

The following procedure outlines how to connect Amazon Q to Amazon WorkDocs using the AWS Management Console.

Connecting Amazon Q to Amazon WorkDocs

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **Amazon WorkDocs** page, enter the following information:
6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. In **Source**, enter the following:
 - **Organization ID specific to your Amazon WorkDocs site** – Select a Amazon WorkDocs directory or create a new one. Only already created directories are available to connect.
 - **Amazon WorkDocs site name** – Enter your Amazon WorkDocs site name.
8. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

 **Note**

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

9. **Configure VPC and security group – optional** – Choose whether you want to use a VPC. If you do, enter the following information:
 - a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
 - b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

10. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

11. In **Sync scope** – Choose what to sync from your data source.

- **Crawl document comments** – Choose to crawl document comments.
 - **regex patterns** – Add regex patterns to include or exclude file names, file types, or file paths. You can have a total of 100 patterns.
12. In **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.
- **Full sync** – Sync all content regardless of the previous sync status.
 - **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.
13. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
14. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
15. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:
- a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
 - b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

 **Note**

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

16. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

17. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

Note

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to Amazon WorkDocs using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the `configuration` parameter to provide a JSON schema with all other configuration information specific to your data source connector.

For an example of the API request, see [CreateDataSource](#) in the Amazon Q API Reference.

Amazon WorkDocs JSON schema

The following is the Amazon WorkDocs JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
      "type": "object",
      "properties": {
        "repositoryEndpointMetadata": {
          "type": "object",
          "properties": {
            "organizationId": {
              "type": "string",
              "minLength": 12,
              "maxLength": 12,
              "pattern": "d-[0-9a-fA-F]{10}"
            },
            "siteName": {
```

```

        "type": "string"
      },
      "domainName": {
        "type": "string"
      }
    },
    "required": ["organizationId"]
  }
}
},
"repositoryConfigurations": {
  "type": "object",
  "properties": {
    "All": {
      "type": "object",
      "properties": {
        "fieldMappings": {
          "type": "array",
          "items": [
            {
              "type": "object",
              "properties": {
                "indexFieldName": {
                  "type": "string"
                },
                "indexFieldType": {
                  "type": "string",
                  "enum": ["STRING", "STRING_LIST", "DATE", "LONG"]
                },
                "dataSourceFieldName": {
                  "type": "string"
                },
                "dateFieldFormat": {
                  "type": "string",
                  "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
                }
              }
            }
          ]
        },
        "required": [
          "indexFieldName",
          "indexFieldType",
          "dataSourceFieldName"
        ]
      }
    }
  ]
}

```



```
    }
  },
  "required": ["fieldMappings"]
},
"required": ["All"]
},
"additionalProperties": {
  "type": "object",
  "properties": {
    "isCrawlAcl": {
      "type": "boolean"
    },
    "fieldForUserId": {
      "type": "string"
    },
    "crawlComments": {
      "type": "string"
    },
    "exclusionPatterns": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    "inclusionPatterns": {
      "type": "array",
      "items": {
        "type": "string"
      }
    }
  }
},
"required": []
},
"enableIdentityCrawler": {
  "type": "boolean"
},
"syncMode": {
  "type": "string",
  "enum": [
    "FORCED_FULL_CRAWL",
    "CHANGE_LOG"
  ]
},
```

```

    "type" : {
      "type" : "string",
      "pattern": "WORKDOCS"
    }
  },
  "version": {
    "type": "string",
    "anyOf": [
      {
        "pattern": "1.0.0"
      }
    ]
  },
  "required": [
    "connectionConfiguration",
    "repositoryConfigurations",
    "syncMode",
    "enableIdentityCrawler",
    "additionalProperties",
    "type"
  ]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.
repositoryEndpointMetadata	The endpoint information for the data source.
organizationId	The identifier of the directory corresponding to your Amazon WorkDocs site repository. You can find the organization ID in the AWS Directory Service by going to Active Directory , then Directories .
siteName	The site of the Amazon WorkDocs site.
domainName	The domain of the Amazon WorkDocs site.

Configuration	Description
<code>repositoryConfigurations</code>	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings.
<ul style="list-style-type: none"> • All 	A list of objects that map the attributes or field names of your Amazon WorkDocs content to Amazon Q index field names.
<code>additionalProperties</code>	Additional configuration options for your content in your data source.
<code>isCrawlAcl</code>	Specify <code>true</code> to crawl ACL information.
<code>fieldForUserId</code>	
<code>crawlComments</code>	Specify <code>true</code> to crawl pages.
<ul style="list-style-type: none"> • <code>exclusionPatterns</code> 	A list of regular expression patterns to exclude specific content from your Amazon WorkDocs data source. Content that matches the patterns are excluded from the index. Content that doesn't match the patterns are excluded from the index. If content matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence, and the content isn't included in the index.
<ul style="list-style-type: none"> • <code>inclusionPatterns</code> 	A list of regular expression patterns to include specific content in your Amazon WorkDocs data source. Content that matches the patterns are included in the index. Content that doesn't match the patterns are excluded in the index. If content matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence, and the content isn't included in the index.

Configuration	Description
<code>type</code>	The type of data source. Specify <code>WORKDOCS</code> as your data source type.
<code>enableIdentityCrawler</code>	Specify <code>true</code> to use the Amazon Q identity crawler to sync identity/principal information on users and groups with access to specific documents. See Identity crawler for more information.
<code>syncMode</code>	Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose between the following options: <ul style="list-style-type: none"> Use <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index. Use <code>CHANGE_LOG</code> to incrementally crawl only new and modified content each time your data source syncs with your index.
<code>version</code>	The version of this template that's currently supported.

How Amazon Q Business connector crawls Amazon WorkDocs ACLs

When you connect an Amazon WorkDocs data source to Amazon Q Business, Amazon Q Business crawls ACL information attached to a document (user and group information) from your Amazon WorkDocs instance. If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

The Amazon WorkDocs group and user IDs are mapped as follows:

- `_group_ids`—Group IDs exist in Amazon WorkDocs on files where there are set access permissions. They are mapped from the names of the groups in Amazon WorkDocs.

- `_user_id`—User IDs exist in Amazon WorkDocs on files where there are set access permissions. They are mapped from the user names in Amazon WorkDocs.

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business Amazon WorkDocs data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q Business enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a mapping already available, or if you want to map additional document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q Amazon WorkDocs connector supports the following entities and the associated reserved and custom attributes.

Amazon WorkDocs field name	Index field name	Description	Data type
id	_document_id	Default	String
authors	_authors	Default	String list
createdTime	_created_at	Default	Date
displayUrl	_source_uri	Default	String
version	_version	Default	String
fileExtension	_file_type	Default	String
category	_category	Default	String
modifiedTime	_last_updated_at	Default	Date

IAM role for Amazon Q Business Amazon WorkDocs connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the BatchPutDocument and BatchDeleteDocument operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.

- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToGetSecret",
      "Effect": "Allow",
      "Action": [
        "secretsmanager:GetSecretValue"
      ],
      "Resource": [
        "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[{{secret_id}}]"
      ]
    },
    {
      "Sid": "AllowsAmazonQToDecryptSecret",
      "Effect": "Allow",
      "Action": [
        "kms:Decrypt"
      ],
      "Resource": [
        "arn:aws:kms:{{region}}:{{account_id}}:key/[{{key_id}}]"
      ],
      "Condition": {
        "StringLike": {
          "kms:ViaService": [
            "secretsmanager.*.amazonaws.com"
          ]
        }
      }
    }
  ],
  {
    "Sid": "AllowsAmazonQToIngestDocuments",
    "Effect": "Allow",
    "Action": [
      "qbusiness:BatchPutDocument",
      "qbusiness:BatchDeleteDocument"
    ]
  }
]
```

```

    "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
    {{application_id}}/index/{{index_id}}"
  },
  {
    "Sid": "AllowsAmazonQToIngestPrincipalMapping",
    "Effect": "Allow",
    "Action": [
      "qbusiness:PutGroup",
      "qbusiness:CreateUser",
      "qbusiness>DeleteGroup",
      "qbusiness:UpdateUser",
      "qbusiness:ListGroups"
    ],
    "Resource": [
      "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}",
      "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
      index/{{index_id}}",
      "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
      index/{{index_id}}/data-source/*"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateAndDeleteNI",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2>DeleteNetworkInterface"
    ],
    "Resource": [
      "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
      "arn:aws:ec2:{{region}}:{{account_id}}:security-group/[{{security_group}}]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2>DeleteNetworkInterface"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:RequestTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
      }
    }
  }
}

```



```

    },
    "ForAllValues:StringEquals": {
      "aws:TagKeys": [
        "AMAZON_Q"
      ]
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateTags",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateTags"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringEquals": {
        "ec2:CreateAction": "CreateNetworkInterface"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterfacePermission"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:ResourceTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
    "Effect": "Allow",
    "Action": [
      "ec2:DescribeNetworkInterfaces",
      "ec2:DescribeAvailabilityZones",
      "ec2:DescribeNetworkInterfaceAttribute",
      "ec2:DescribeVpcs",
      "ec2:DescribeRegions",
      "ec2:DescribeNetworkInterfacePermissions",

```

```

    "ec2:DescribeSubnets"
  ],
  "Resource": "*"
}
]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "{{source_account}}"
        },
        "ArnEquals": {
          "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
        }
      }
    }
  ]
}

```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Troubleshooting your Amazon Q Business Amazon WorkDocs connector

The following table provides information about error codes you may see for the Amazon WorkDocs connector and suggested troubleshooting actions.

Error code	Error message	Suggested resolution
5001	Add valid organizationId.	organizationId should not be null or empty.
5002	Add valid domainName.	domainName should not be null or empty.
5003	Add valid siteName.	SiteName should not be null or empty.
5004	There was an error parsing the field value.	Size has exceeded the maximum allowable limit.
5005	Error in de-serializing change log token.	Wait for a few minutes and try again.
5006	Error in serializing change log token.	Wait for a few minutes and try again.
5007	Amazon WorkDocs Service is not available.	Wait for a few minutes and try again.
5008	Amazon Q is unable to reach Amazon WorkDocs Server at this moment.	Wait for a few minutes and try again.
5009	Amazon Q is unable to assume index IAM role.	Ensure that service principal qbusiness.amazonaws.com is added to IAM role trust policy.
5010	Operation is not permitted.	Wait for a few minutes and try again.
5100	There was a problem while retrieving the values for the field.	The values may be empty or incorrect. It should be either true or false.

Error code	Error message	Suggested resolution
5099	An exception has occurred while calling Amazon WorkDocs API.	Add permissions to call API in your data source IAM role.
5098	Amazon Q is unable to find id.	Wait for a few minutes and try again.

Connecting Box to Amazon Q Business

Box is a cloud storage service that offers file hosting capabilities. You can connect your Box instance to Amazon Q Business—using either the AWS Management Console, CLI, or the [CreateDataSource](#) API—and create an Amazon Q web experience.

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).
- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [Box connector overview](#)
- [Prerequisites for connecting Amazon Q Business to Box](#)
- [Connecting Amazon Q Business to Box using the console](#)
- [Connecting Amazon Q Business to Box using APIs](#)
- [How Amazon Q Business connector crawls Box ACLs](#)
- [Amazon Q Business Box data source connector field mappings](#)
- [IAM role for Amazon Q Business Box connector](#)
- [Known limitations for the Amazon Q Box connector](#)

Box connector overview

The following table gives an overview of the Amazon Q Business Box connector and its supported features.

Category	Feature	Support
Security	Authentication type	OAuth 2.0
	Authentication credentials	<ul style="list-style-type: none"> Client ID Client secret Public Key ID Private Key Pass Phrase <div style="border: 1px solid #f08080; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p> Important Admin privileges required.</p> </div>
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Identity crawling	Yes
	VPC	Yes
Crawl features	Custom metadata	Yes
	Entities	<p>Yes. The following entities are supported:</p> <ul style="list-style-type: none"> Files Comments Tasks Web links

Category	Feature	Support
	Field mappings	Yes. Supports both default and custom field mappings. For more information, see Field mappings .
	Filters	Yes. The following filters are supported: <ul style="list-style-type: none"> • Include web links • Include comments • Include tasks • Include/exclude by file name • Include/exclude by file type • Include/exclude by file path
	Sync mode	Supports full and incremental sync.
	File types	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to Box

Before you begin, make sure that you have completed the following prerequisites.

In Box, make sure you have:

- A Box Enterprise or Box Enterprise Plus account.
- Created a Box custom app in the Box Developer Console and configured it to use **Server Authentication (with JWT)**.
- Set your **App Access Level** to **App + Enterprise Access** and allowed it to **Make API calls using the as-user header**.
- Used the admin user to add the following **Application Scopes** in your Box app:
 - Write all files and folders stored in a Box
 - Manage users
 - Manage groups
 - Manage enterprise properties

- Generated and downloaded Public/Private key pair including a client ID, a client secret, a public key ID, private key ID, a pass phrase, and an enterprise ID to use as authentication credentials. See [Public and private keypair](#) for more details.
- Copied your Box enterprise ID either from your Box Developer Console settings or from your Box app. For example, *801234567*.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your Box authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Connecting Amazon Q Business to Box using the console

The following procedure outlines how to connect Amazon Q Business to Box using the AWS Management Console.

Connecting Amazon Q to Box

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **Box** page, enter the following information:
6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.
7. **Source** – Enter your **Box enterprise ID**.
8. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

 **Note**

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

9. In **Authentication** – Choose to create an **AWS Secrets Manager secret** and then enter the following information for your **AWS Secrets Manager secret**.
 - a. **Secret name** – A name for your secret.
 - b. **Client ID** – The client ID provided by Box.
 - c. **Client Secret** – The client secret provided by Box.
 - d. **Public Key ID** – Your Box public key ID.
 - e. **Private Key** – The private key provided by Box.
 - f. **Pass Phrase** – The pass phrase you use to log into your Box account.
10. **Identity crawler** – Choose to activate Amazon Q identity crawler to sync identity information. For more information, see [Identity crawler](#).
11. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

12. **Configure VPC and security group** – *optional* – Choose whether you want to use a VPC. If you do, enter the following information:
 - a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.

- b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

13. In **Sync scope**, enter the following information:

- a. **Select additional kinds of content to index** – Choose whether to include **Web links**, **Comments**, and **Tasks**.

 **Note**

Box files are indexed by default.

- b. **Additional configuration – optional** – Configure the following settings:

- **Regex patterns** – Regular expression patterns to include or exclude certain files. You can add up to 100 patterns.

14. For **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.

- **Full sync** – Sync all content regardless of the previous sync status.
- **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.

15. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).

16. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.

17. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:

- a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
- b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

Note

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

18. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

19. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

Note

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to Box using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the `configuration` parameter to provide a JSON schema with all other configuration information specific to your data source connector.

For an example of the API request, see [CreateDataSource](#) in the Amazon Q API Reference.

Box JSON schema

The following is the Box JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
```

```
"properties": {
  "connectionConfiguration": {
    "type": "object",
    "properties": {
      "repositoryEndpointMetadata": {
        "type": "object",
        "properties": {
          "enterpriseId": {
            "type": "string",
            "minLength": 9,
            "maxLength": 9
          }
        },
        "required": [
          "enterpriseId"
        ]
      },
      "required": [
        "repositoryEndpointMetadata"
      ]
    },
    "required": [
      "repositoryEndpointMetadata"
    ]
  },
  "repositoryConfigurations": {
    "type": "object",
    "properties": {
      "file": {
        "type": "object",
        "properties": {
          "fieldMappings": {
            "type": "array",
            "items": [
              {
                "type": "object",
                "properties": {
                  "indexFieldName": {
                    "type": "string"
                  },
                  "indexFieldType": {
                    "type": "string",
                    "enum": [
                      "STRING",
                      "STRING_LIST",
                      "DATE",
                      "LONG"
                    ]
                  }
                }
              ]
            ]
          }
        }
      }
    }
  }
}
```

```
    ]
    },
    "dataSourceFieldName": {
      "type": "string"
    },
    "dateFieldFormat": {
      "type": "string",
      "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
    }
  },
  "required": [
    "indexFieldName",
    "indexFieldType",
    "dataSourceFieldName"
  ]
}
]
}
},
"required": [
  "fieldMappings"
]
},
"task": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": [
        {
          "type": "object",
          "properties": {
            "indexFieldName": {
              "type": "string"
            },
            "indexFieldType": {
              "type": "string",
              "enum": [
                "STRING",
                "STRING_LIST",
                "DATE",
                "LONG"
              ]
            }
          }
        }
      ]
    }
  }
},
```

```

        "dataSourceFieldName": {
            "type": "string"
        },
        "dateFieldFormat": {
            "type": "string",
            "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
        }
    },
    "required": [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
    ]
}
]
}
},
"required": [
    "fieldMappings"
]
},
"comment": {
    "type": "object",
    "properties": {
        "fieldMappings": {
            "type": "array",
            "items": [
                {
                    "type": "object",
                    "properties": {
                        "indexFieldName": {
                            "type": "string"
                        },
                        "indexFieldType": {
                            "type": "string",
                            "enum": [
                                "STRING",
                                "STRING_LIST",
                                "DATE",
                                "LONG"
                            ]
                        }
                    }
                }
            ]
        },
        "dataSourceFieldName": {
            "type": "string"
        }
    }
}

```

```

        },
        "dateFieldFormat": {
            "type": "string",
            "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
        }
    },
    "required": [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
    ]
}
]
}
},
"required": [
    "fieldMappings"
]
},
"webLink": {
    "type": "object",
    "properties": {
        "fieldMappings": {
            "type": "array",
            "items": [
                {
                    "type": "object",
                    "properties": {
                        "indexFieldName": {
                            "type": "string"
                        },
                        "indexFieldType": {
                            "type": "string",
                            "enum": [
                                "STRING",
                                "STRING_LIST",
                                "DATE",
                                "LONG"
                            ]
                        },
                        "dataSourceFieldName": {
                            "type": "string"
                        },
                        "dateFieldFormat": {

```

```
        "type": "string",
        "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
      }
    },
    "required": [
      "indexFieldName",
      "indexFieldType",
      "dataSourceFieldName"
    ]
  }
]
}
},
"required": [
  "fieldMappings"
]
}
}
},
"additionalProperties": {
  "type": "object",
  "properties": {
    "isCrawlAcl": {
      "type": "boolean"
    },
    "fieldForUserId": {
      "type": "string"
    },
    "crawlComments": {
      "type": "boolean"
    },
    "crawlTasks": {
      "type": "boolean"
    },
    "crawlWebLinks": {
      "type": "boolean"
    },
    "inclusionPatterns": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    "exclusionPatterns": {
```

```
        "type": "array",
        "items": {
          "type": "string"
        }
      },
      "required": []
    },
    "type": {
      "type": "string",
      "pattern": "BOX"
    },
    "enableIdentityCrawler": {
      "type": "boolean"
    },
    "syncMode": {
      "type": "string",
      "enum": [
        "FULL_CRAWL",
        "FORCED_FULL_CRAWL",
        "CHANGE_LOG"
      ]
    },
    "secretArn": {
      "type": "string",
      "minLength": 20,
      "maxLength": 2048
    }
  },
  "version": {
    "type": "string",
    "anyOf": [
      {
        "pattern": "1.0.0"
      }
    ]
  }
},
"required": [
  "connectionConfiguration",
  "repositoryConfigurations",
  "syncMode",
  "additionalProperties",
  "secretArn",
  "type",
```



```

    "enableIdentityCrawler"
  ]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.
repositoryEndpointMetadata	The endpoint information for the data source.
enterpriseId	The Box enterprise id.
repositoryConfigurations	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings.
<ul style="list-style-type: none"> • file • task • comment • webLink 	A list of objects that map the attributes or field names of your Box files, tasks, comments, and webLinks to Amazon Q index field names.
additionalProperties	Additional configuration options for your content in your data source.
isCrawlAcl	Specify true to crawl access control information from documents.
fieldForUserId	Specify field to use for UserId for ACL crawling.
crawlComments	Specify true to crawl assets.
crawlTasks	Specify true to crawl pages.
crawlWebLinks	Specify true to crawl pages.

Configuration	Description
<ul style="list-style-type: none"> • InclusionPatterns • ExclusionPatterns 	<p>A list of regular expression patterns to include or exclude specific content from your Box data source. Content that matches the patterns are included in the index. Content that doesn't match the patterns are excluded from the index. If content matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence, and the content isn't included in the index.</p>
type	<p>The type of data source. Specify BOX as your data source type.</p>
enableIdentityCrawler	<p>Specify true to use the Amazon Q identity crawler to sync identity/principal information on users and groups with access to specific documents.</p>
syncMode	<p>Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose between the following options:</p> <ul style="list-style-type: none"> • Use <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index. • Use <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index. • Use <code>CHANGE_LOG</code> to incrementally crawl only new and modified content each time your data source syncs with your index.

Configuration	Description
secretArn	<p>The Amazon Resource Name (ARN) of an AWS Secrets Manager secret that contains the key-value pairs required to connect to your Box. The secret must contain a JSON structure with the following keys:</p> <pre data-bbox="829 489 1507 806"> { "clientId": " <i>client-id</i> ", "clientSecret": " <i>client-secret</i> ", "publicKeyId": " <i>public-key-id</i> ", "privateKey": " <i>private-key</i> ", "passphrase": " <i>pass-phrase</i> " } </pre>
version	The version of this template that's currently supported.

How Amazon Q Business connector crawls Box ACLs

When you connect an Box data source to Amazon Q Business, Amazon Q Business crawls ACL information attached to a document (user and group information) from your Box instance. If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

- `_group_ids`—Group IDs exist in Box on files where there are set access permissions. They are mapped from the names of the groups in Box.
- `_user_id`—User IDs exist in Box on files where there are set access permissions. They are mapped from the user emails as the user IDs in Box.

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business Box data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q Business enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a attribute mapping already available, or if you want to map additional document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q Box connector supports the following entities and the associated reserved and custom attributes.

Supported entities and field mappings

- [Files and folders](#)
- [Comments](#)
- [Tasks](#)
- [Web links](#)

Files and folders

Box field name	Index field name	Description	Data type
bx_createdAt	_created_at	Default	Date
bx_modifiedAt	_last_updated_at	Default	Date
bx_authors	_authors	Default	String list
bx_uri	_source_uri	Default	String
bx_size	bx_file_size	Custom	String
bx_category	_category	Default	String

Comments

Box field name	Index field name	Description	Data type
bx_createdAt	_created_at	Default	Date
bx_modifiedAt	_last_updated_at	Default	Date
bx_author	_authors	Custom	String
bx_parentFile	bx_comment_item	Custom	String
bx_category	_category	Default	String

Tasks

Box field name	Index field name	Description	Data type
bx_createdAt	_created_at	Default	Date
bx_action	bx_task_action	Custom	String
bx_taskComplete	bx_task_completed	Custom	String

Box field name	Index field name	Description	Data type
bx_taskItem	bx_task_item	Custom	String
bx_taskAssigned	bx_task_assigned_to	Custom	String
bx_author	bx_author	Custom	String
bx_category	_category	Default	String
bx_uri	_source_uri	Default	String

Web links

Box field name	Index field name	Description	Data type
bx_createdAt	_created_at	Default	Date
bx_author	bx_author	Custom	String
bx_category	_category	Default	String
bx_uri	_source_uri	Default	String

IAM role for Amazon Q Business Box connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the BatchPutDocument and BatchDeleteDocument operations to ingest documents.

- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToGetSecret",
      "Effect": "Allow",
      "Action": [
        "secretsmanager:GetSecretValue"
      ],
      "Resource": [
        "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
      ]
    },
    {
      "Sid": "AllowsAmazonQToDecryptSecret",
      "Effect": "Allow",
      "Action": [
        "kms:Decrypt"
      ],
      "Resource": [
        "arn:aws:kms:{{region}}:{{account_id}}:key/[key_id]"
      ],
      "Condition": {
        "StringLike": {
          "kms:ViaService": [
            "secretsmanager.*.amazonaws.com"
          ]
        }
      }
    },
    {
      "Sid": "AllowsAmazonQToIngestDocuments",
      "Effect": "Allow",
      "Action": [
        "qbusiness:BatchPutDocument",
```

```

    "qbusiness:BatchDeleteDocument"
  ],
  "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
},
{
  "Sid": "AllowsAmazonQToIngestPrincipalMapping",
  "Effect": "Allow",
  "Action": [
    "qbusiness:PutGroup",
    "qbusiness:CreateUser",
    "qbusiness>DeleteGroup",
    "qbusiness:UpdateUser",
    "qbusiness:ListGroup"
  ],
  "Resource": [
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}",
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}",
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}/data-source/*"
  ]
},
{
  "Sid": "AllowsAmazonQToCreateAndDeleteNI",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateNetworkInterface",
    "ec2>DeleteNetworkInterface"
  ],
  "Resource": [
    "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
    "arn:aws:ec2:{{region}}:{{account_id}}:security-group/[{{security_group}}]"
  ]
},
{
  "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateNetworkInterface",
    "ec2>DeleteNetworkInterface"
  ],
  "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
  "Condition": {

```



```

    "StringLike": {
      "aws:RequestTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
    },
    "ForAllValues:StringEquals": {
      "aws:TagKeys": [
        "AMAZON_Q"
      ]
    }
  }
},
{
  "Sid": "AllowsAmazonQToCreateTags",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateTags"
  ],
  "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
  "Condition": {
    "StringEquals": {
      "ec2:CreateAction": "CreateNetworkInterface"
    }
  }
},
{
  "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateNetworkInterfacePermission"
  ],
  "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
  "Condition": {
    "StringLike": {
      "aws:ResourceTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
    }
  }
},
{
  "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
  "Effect": "Allow",
  "Action": [
    "ec2:DescribeNetworkInterfaces",
    "ec2:DescribeAvailabilityZones",
    "ec2:DescribeNetworkInterfaceAttribute",
    "ec2:DescribeVpcs",

```

```

        "ec2:DescribeRegions",
        "ec2:DescribeNetworkInterfacePermissions",
        "ec2:DescribeSubnets"
    ],
    "Resource": "*"
}
]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "{{source_account}}"
        },
        "ArnEquals": {
          "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
        }
      }
    }
  ]
}

```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Known limitations for the Amazon Q Box connector

The Amazon Q Box connector has the following known limitations:

- Crawling data from external folders is not supported.

Connecting Confluence (Cloud) to Amazon Q Business

Atlassian Confluence is a collaborative work-management tool designed for sharing, storing, and working on project planning, software development, and product management. You can connect Confluence (Cloud) instance to Amazon Q Business—using either the AWS Management Console or the [CreateDataSource](#) API—and create an Amazon Q web experience.

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).
- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [Confluence \(Cloud\) connector overview](#)
- [Prerequisites for connecting Amazon Q Business to Confluence \(Cloud\)](#)
- [Setting up Confluence \(Cloud\) for connecting to Amazon Q Business](#)
- [Connecting Amazon Q Business to Confluence \(Cloud\) using the console](#)
- [Connecting Amazon Q Business to Confluence \(Cloud\) using APIs](#)
- [How Amazon Q Business connector crawls Confluence \(Cloud\) ACLs](#)
- [Amazon Q Business Confluence \(Cloud\) data source connector field mappings](#)
- [IAM role for Amazon Q Confluence \(Cloud\) connector](#)
- [Troubleshooting your Amazon Q Business Confluence \(Cloud\) connector](#)

Confluence (Cloud) connector overview

The following table gives an overview of the Amazon Q Business Confluence (Cloud) connector and its supported features.

Category	Feature	Support
Security	Authentication type	Basic, OAuth 2.0

Category	Feature	Support
	Authentication credentials	<p>For Basic authentication</p> <ul style="list-style-type: none"> • Confluence Cloud URL • Confluence username • Password (Confluence (Cloud) site token) <p>For OAuth 2.0 authentication</p> <ul style="list-style-type: none"> • App key • App secret • Access token • Refresh token <div style="border: 1px solid #add8e6; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p> Note</p> <p>Access and refresh tokens expire in 1 hour. For information on regenerating tokens, see Atlassian Developer Documentation.</p> </div>
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Identity crawling	Yes
Crawl features	Custom metadata	Yes
	Entities	<p>Yes. The following entities are supported:</p> <ul style="list-style-type: none"> • Space • Page • Blog post • Comment • Attachment

Category	Feature	Support
	Field mappings	Yes. For more information, see Field mappings .
	Filters	Yes. The following filters are supported: <ul style="list-style-type: none"> • Inclusion exclusion filters for Space key and Space URL • Inclusion exclusion filters on File Type for Attachment entity • Supports regex filters for entities • Supports inclusion and exclusion filters for File size
	Sync mode	Supports full and incremental (new, modified, and deleted) sync.
	File types	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to Confluence (Cloud)

Before you begin, make sure that you have completed the following prerequisites.

In Confluence Cloud, make sure you have:

- Copied your Confluence instance URL. For example: <https://example.atlassian.net>. You need your Confluence instance URL to connect to Amazon Q.
- Configured basic authentication credentials containing a username (email ID used to log into Confluence) and password (Confluence API token) to allow Amazon Q to connect to your Confluence instance. For information about how to create a Confluence API token, see [Manage API tokens for your Atlassian account](#) on the Atlassian website.
- **Optional:** Configured OAuth 2.0 credentials containing a Confluence app key, Confluence app secret, Confluence access token, and Confluence refresh token to allow Amazon Q to connect to your Confluence instance. If your access token expires, you can either use the refresh token to regenerate your access token and refresh token pair. Or, you can repeat the authorization process. For more information about access tokens, see [Manage OAuth access tokens](#) on the Atlassian website.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your Confluence (Cloud) authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Setting up Confluence (Cloud) for connecting to Amazon Q Business

Before you connect Confluence (Cloud) to Amazon Q Business, you need to create and retrieve the Confluence (Cloud) credentials you will use to connect Confluence (Cloud) to Amazon Q. You will also need to add any permissions needed by Confluence (Cloud) to connect to Amazon Q.

The following procedure gives you an overview of how to configure Confluence (Cloud) to connect to Amazon Q using either basic authentication or OAuth 2.0 authentication.

Topics

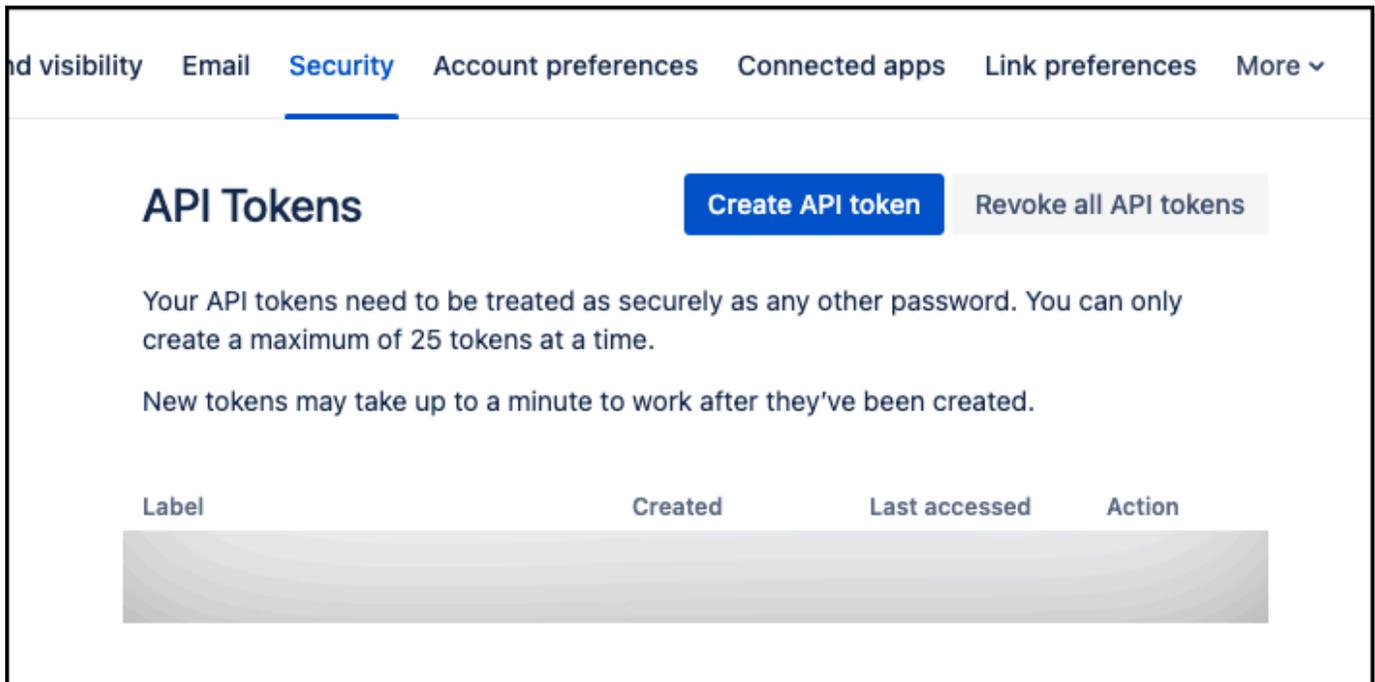
- [Basic authentication](#)
- [OAuth 2.0 authentication](#)
- [How Amazon Q works with Confluence \(Cloud\) access and refresh tokens](#)

Basic authentication

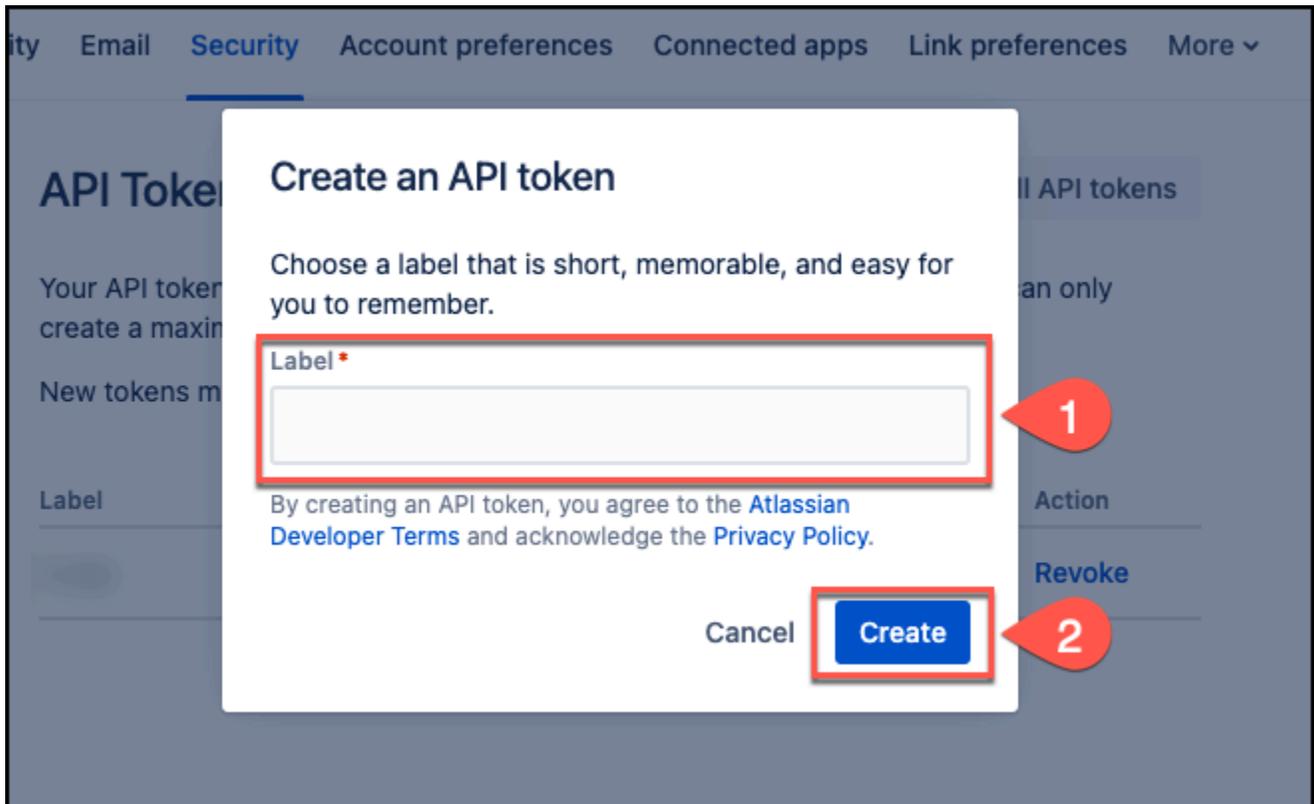
You can connect Amazon Q to Confluence (Cloud) using basic authentication credentials. The following procedure gives you an overview of how to configure Confluence (Cloud) to connect to Amazon Q using basic authentication.

Configuring Confluence (Cloud) basic authentication for Amazon Q

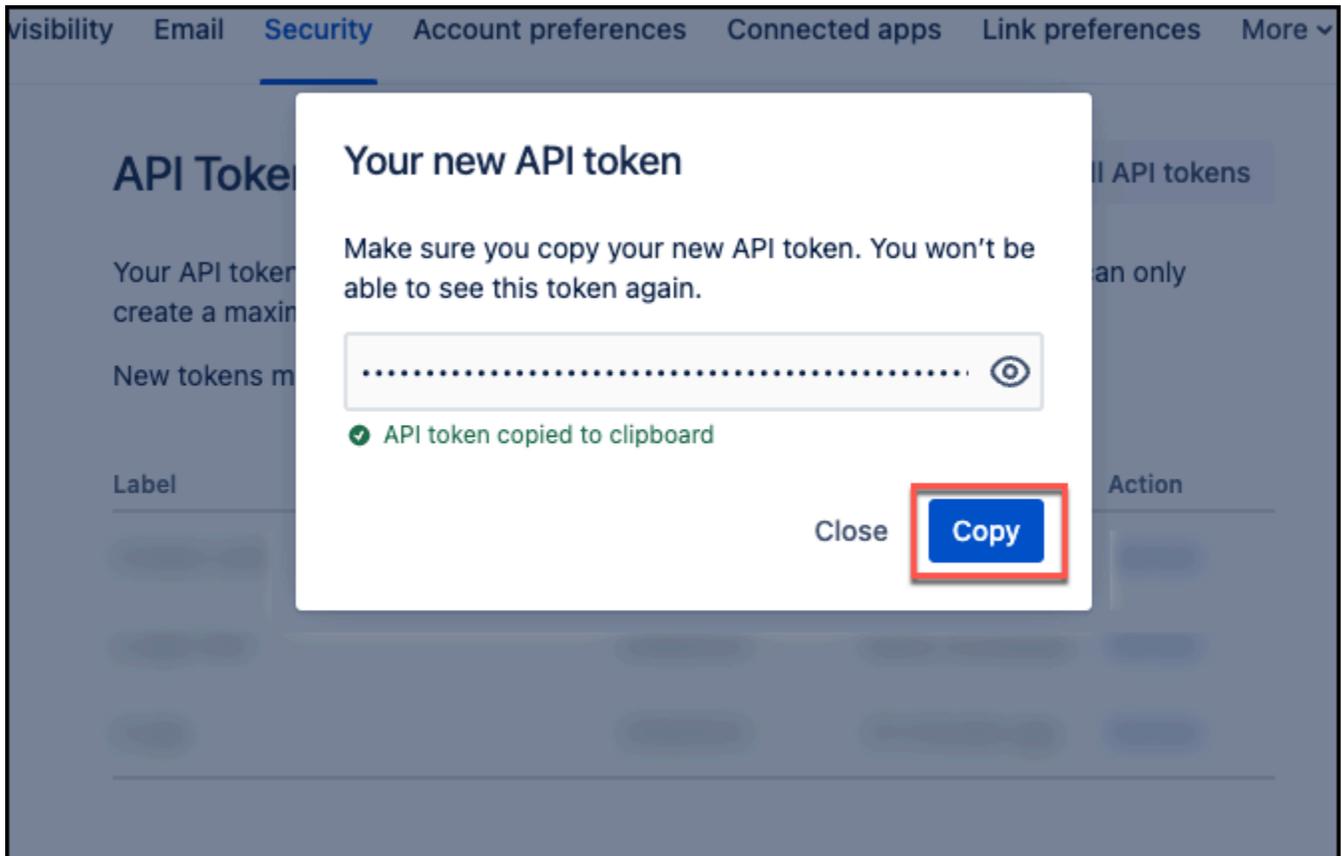
1. Log in to your account from the [Confluence \(Cloud\)](#). Note the username you logged in with. You will need this later to connect to Amazon Q.
2. From your Confluence (Cloud) home page, note your Confluence (Cloud) URL from your Confluence browser URL. For example: *https://example.atlassian.net*. You will need this later to connect to Amazon Q.
3. Then, go to [Security](#) page in Confluence (Cloud).
4. From the **API tokens** page, select **Create API token**.



5. In the **Create an API token** dialog box that opens, for **Label**, add a name for your API token. Then, select **Create**.



- From the **Your new API token** dialog box, copy the API token and save it in a text editor of your choice. You can't retrieve the API token once you close the dialog box.



7. Select **Close**.

You now have the username, Confluence (Cloud) URL, and Confluence (Cloud) API token you need to connect to Amazon Q with basic authentication.

For more information, see [Manage API tokens for your Atlassian account](#) in Atlassian Support.

OAuth 2.0 authentication

You can connect Amazon Q to Confluence (Cloud) using OAuth 2.0 authentication credentials. The following procedures give you an overview of how to configure Confluence (Cloud) to connect to Amazon Q using OAuth 2.0 authentication.

Steps to configure Confluence (Cloud) OAuth 2.0 authentication

- [Step 1: Retrieving username and Confluence \(Cloud\) URL](#)
- [Step 2: Configuring an OAuth 2.0 app integration](#)
- [Step 3: Retrieving Confluence \(Cloud\) client ID and client Secret](#)
- [Step 4: Generating an Confluence \(Cloud\) access token](#)

- [Step 5: Generating a Confluence \(Cloud\) refresh token](#)
- [Step 6: Generating a new Confluence \(Cloud\) access token using a refresh token](#)

Step 1: Retrieving username and Confluence (Cloud) URL

To connect Confluence (Cloud) to Amazon Q, you need your Confluence (Cloud) username and your Confluence (Cloud) URL. The following procedure shows you how to retrieve these.

Retrieving username and Confluence (Cloud) URL

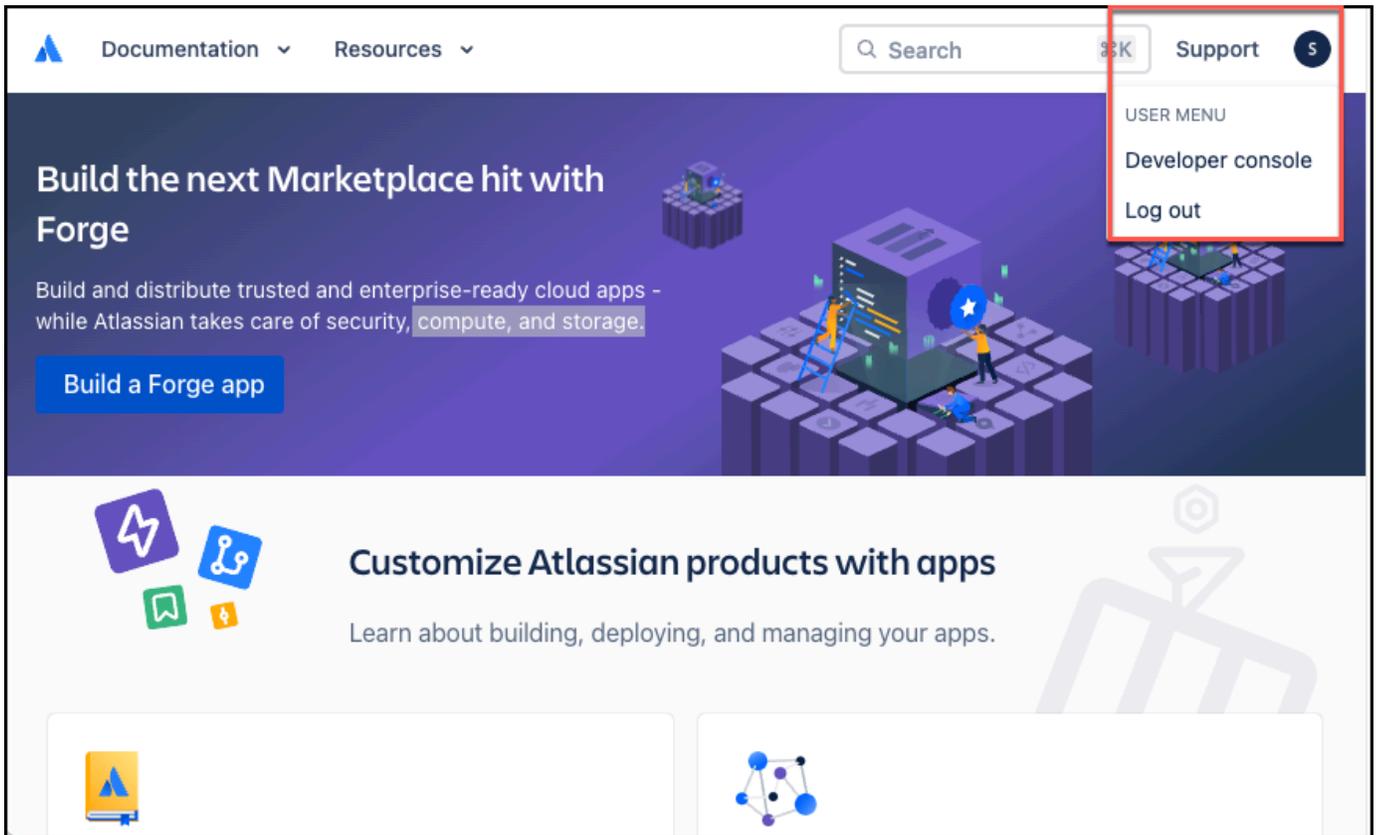
1. Log in to your account from the [Confluence \(Cloud\)](#). Note the username you logged in with. You will need this later to connect to Amazon Q.
2. From your Confluence (Cloud) home page, note your Confluence (Cloud) URL from your Confluence browser URL. For example: *https://example.atlassian.net*. You will need this later to both configure your OAuth 2.0 token and connect to Amazon Q.

Step 2: Configuring an OAuth 2.0 app integration

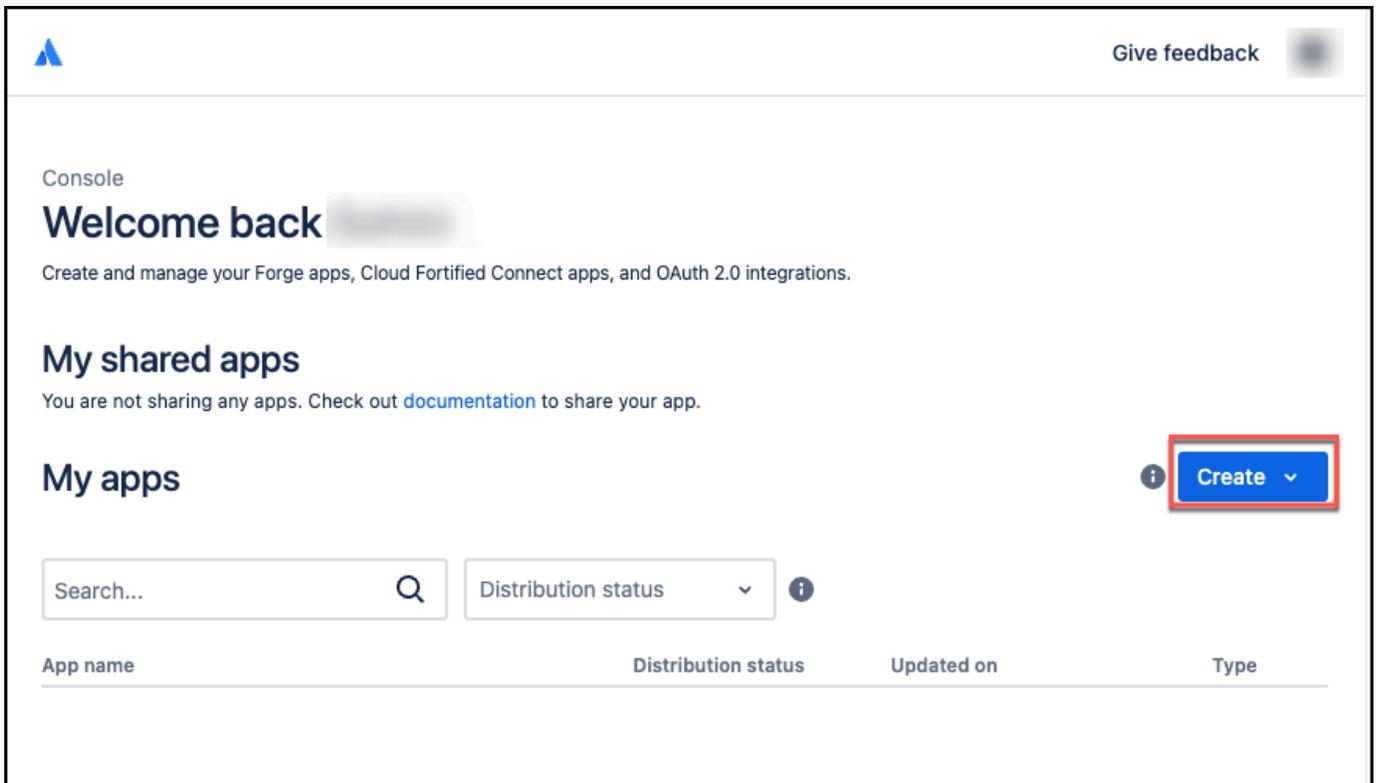
To connect Confluence (Cloud) to Amazon Q using OAuth 2.0 authentication, you need to create a Confluence (Cloud) OAuth 2.0 app with the necessary permissions. The following procedure shows you how to create this.

Configuring an OAuth 2.0 app integration

1. Log in to your account from the [Atlassian Developer page](#).
2. Select the profile icon from the top-right corner. Then, from the dropdown menu that opens, select **Developer Console**.



3. From the **Welcome** page, select **Create** and then select **OAuth 2.0 integration**.



4. On the **Create a new OAuth 2.0 (3LO) integration** page, for **Name**, enter a name for the OAuth 2.0 application you are creating. Then, select the **I agree to be bound by Atlassian's developer terms** checkbox, and select **Create**.

i **Rotating refresh tokens are enabled**

New OAuth 2.0 integrations must use rotating refresh tokens. Rotating refresh tokens improve security by limiting the validity of the refresh token and enabling automatic detection of refresh token reuse.

[Learn more](#) · [Dismiss](#)

Create a new OAuth 2.0 (3LO) integration

An app provides API credentials for Atlassian products and services, as well as features such as OAuth 2.0 (3LO).

1 **Name ***

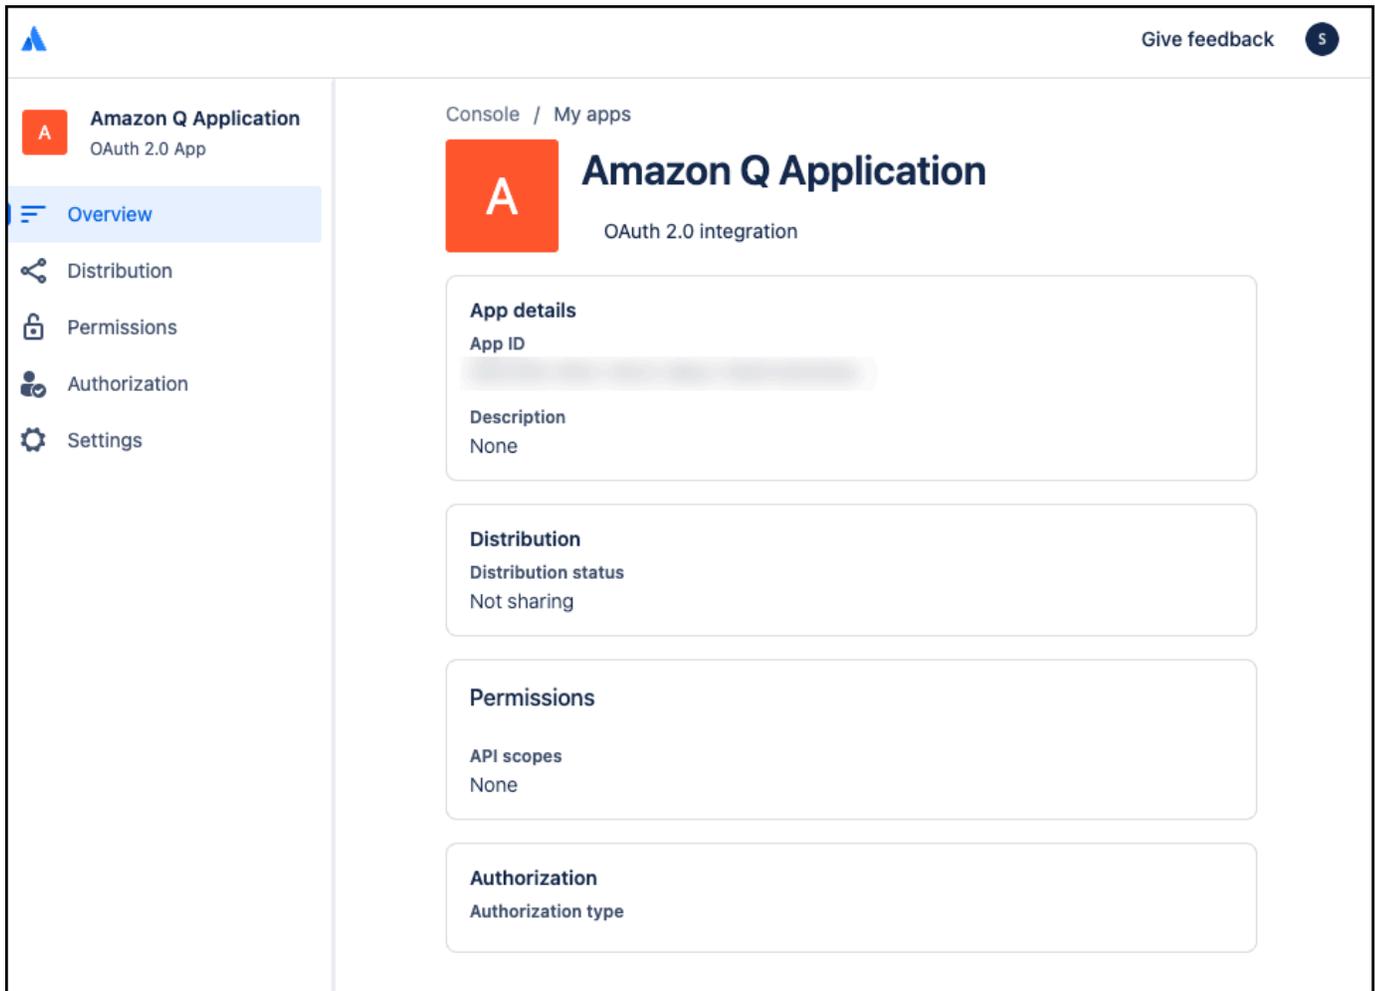
App name

Name your app according to its purpose, for example, Dropbox integration or Timesheets for Jira.

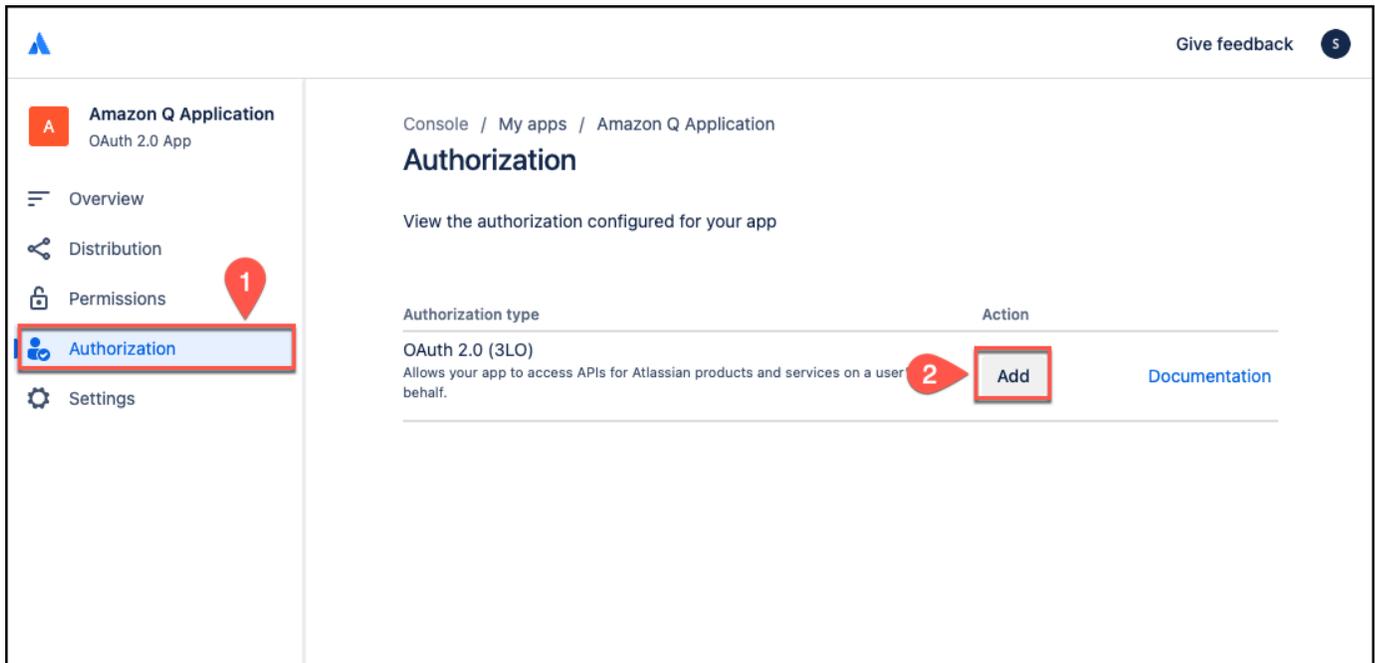
2 I agree to be bound by [Atlassian's developer terms](#).

3

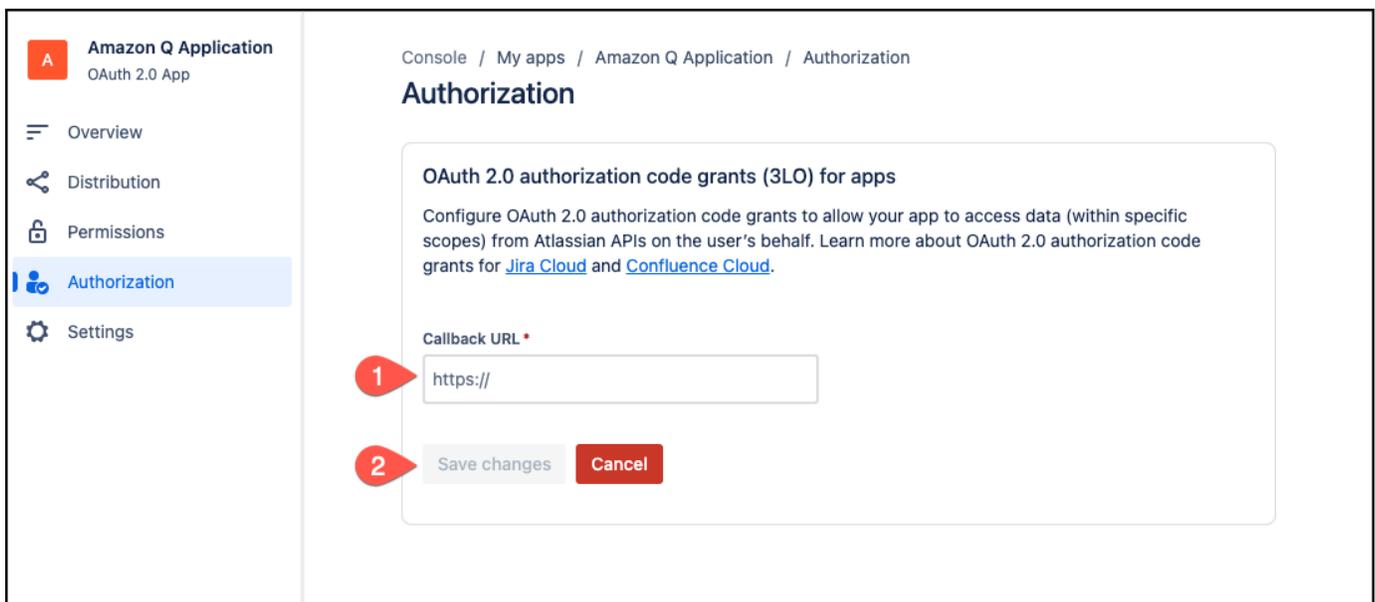
The console will display a summary page outlining the details of the OAuth 2.0 app created.



5. From the left navigation menu, choose **Authorization**.
6. From the **Authorization** page, choose **Add** to add **OAuth 2.0 (3LO)** to your app.



- On the **OAuth 2.0 authorization code grants (3LO) for apps**, enter the Confluence (Cloud) URL you copied as the **Callback URL** and then choose **Save changes**.



- From the **Authorization URL generator** section that appears, choose **Add APIs** to add APIs to your app. This will redirect you to the **Permissions** page.
- On the **Permissions** page, for **Scopes**, navigate to **User Identity API**. Select **Add**, and then select **Configure**.

Permissions

Add and configure your app's API scopes. See [OAuth 2.0 \(3LO\) for apps](#).

Scopes Used
0

Scopes
We recommend that you don't add more than 50 scopes to your app. Use classic scopes to minimize the number of scopes you need. [Learn more](#)

API name	Scopes used	Action
User identity API Get the profile details for the currently logged-in user, such as the Atlassian account ID and email.	0	Add Documentation
Confluence API Get, create, update, and delete content, spaces, and more.	0	Add Documentation
BRIE API Create, cancel, and read backup and restore, retrieve and publish cloud details.	0	Add Documentation
Jira API Get, create, update, and delete issues, projects, fields, and more.	0	Add Documentation
Personal data reporting API Report user accounts that an app is storing personal data for.	0	Add Documentation

10. On the **User Identity API** page, choose **Edit Scopes**, and then add the following read scopes:

- **read:me** – View active user profile
- **read:account** – View user profiles

User identity API

Add and configure your app's API scopes. See [OAuth 2.0 \(3LO\) for apps](#).

Scopes Used
1

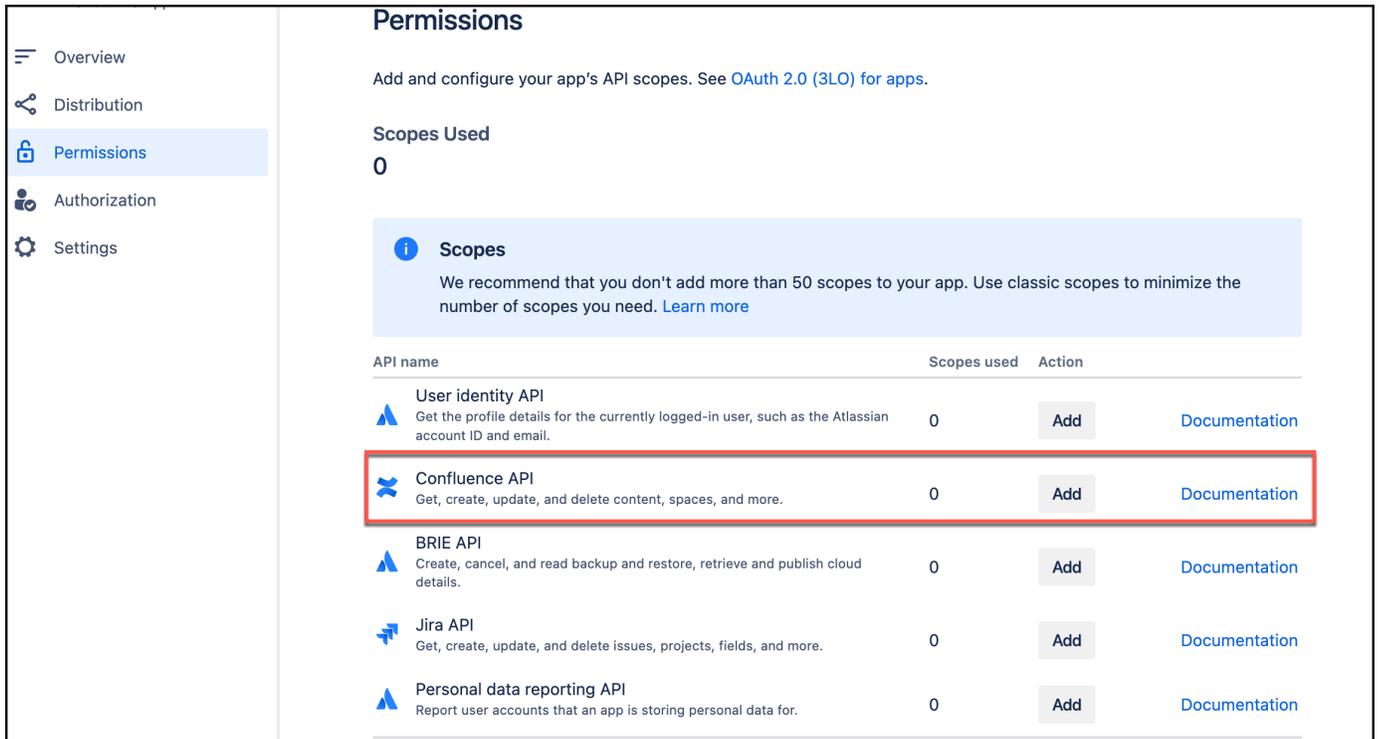
Choosing your scopes
Use the scopes recommended in the API documentation for the features you're using. [Learn more](#)

[Edit Scopes](#)

Select	Scope Name	Code
<input checked="" type="checkbox"/>	View active user profile View the profile details for the currently logged-in user.	read:me
<input type="checkbox"/>	View user profiles Required to view users profiles	read:account

Then, select **Save**.

11. Return to the **Permissions** page. From **Scopes**, navigate to **Confluence API**. Select **Add**, and then select **Configure**.



The screenshot shows the 'Permissions' page in Amazon Q Business. On the left is a navigation menu with 'Permissions' selected. The main content area is titled 'Permissions' and includes a sub-header 'Scopes Used' with a count of '0'. Below this is an informational box about adding scopes. A table lists several API scopes, with the 'Confluence API' row highlighted by a red border. Each row includes an API name, a description, the number of scopes used (all are 0), an 'Add' button, and a 'Documentation' link.

API name	Scopes used	Action
 User identity API Get the profile details for the currently logged-in user, such as the Atlassian account ID and email.	0	Add Documentation
 Confluence API Get, create, update, and delete content, spaces, and more.	0	Add Documentation
 BRIE API Create, cancel, and read backup and restore, retrieve and publish cloud details.	0	Add Documentation
 Jira API Get, create, update, and delete issues, projects, fields, and more.	0	Add Documentation
 Personal data reporting API Report user accounts that an app is storing personal data for.	0	Add Documentation

12. On the **Confluence API** page, make sure you're in the **Classic scopes** section.

Console / My apps / ff / Permissions

Confluence API

Add and configure your app's API scopes. See [OAuth 2.0 \(3LO\) for apps](#).

Scopes Used

0



Choosing your scopes

Use the scopes recommended in the API documentation for the features you're using.

[Learn more](#)

Classic scopes

Granular scopes

Edit Scopes

Select **Scope Name** ⇅

Code ⇅

<input type="checkbox"/>	Write Confluence content Permits the creation of pages, blogs, comments and questions.	<code>write:confluence-content</code>
<input type="checkbox"/>	Read Confluence space summary Read a summary of space information without expansions.	<code>read:confluence-space.summary</code>
<input type="checkbox"/>	Manage Confluence space details Create, update and delete space information.	<code>write:confluence-space</code>
<input type="checkbox"/>	Upload Confluence attachments Upload attachments.	<code>write:confluence-file</code>

Then, choose **Edit Scopes**, and then add the following read scopes:

- **read:confluence-space.summary** – Read Confluence space summary
- **read:confluence-props** – Read Confluence content properties
- **read:confluence-content.all** – Read Confluence detailed content
- **read:confluence-content.summary** – Read Confluence content summary
- **read:confluence-content.permission** – Read content permission in Confluence
- **read:confluence-user** – Read user

- **read:confluence-groups** – Read user groups

Then, select **Save**.

13. Navigate to the **Granular scopes** page.

Console / My apps / ff / Permissions

Confluence API

Add and configure your app's API scopes. See [OAuth 2.0 \(3LO\) for apps](#).

Scopes Used

0

i **Choosing your scopes**

Use the scopes recommended in the API documentation for the features you're using.

[Learn more](#)

Classic scopes **Granular scopes**

Search by name or code All operations All entities

Hide unselected scopes Edit Scopes

Select	Scope Name	Code
<input type="checkbox"/>	View detailed contents View all contents, such as pages, blogposts, whiteboards, comments and attachments.	read:content:confluence
<input type="checkbox"/>	Create and update contents Create and update content, such as pages, blogposts, whiteboards.	write:content:confluence
<input type="checkbox"/>	View content details View details regarding content and its associated properties	read:content-details:confluence

Then, choose **Edit Scopes**, and then add the following read scopes:

- **read:content:confluence** – View detailed contents
- **read:content-details:confluence** – View content details

- **read:space-details:confluence** – View space details
- **read:audit-log:confluence** – View audit records
- **read:page:confluence** – View pages
- **read:attachment:confluence** – View and download content attachments
- **read:blogpost:confluence** – View blogposts
- **read:custom-content:confluence** – View custom content
- **read:comment:confluence** – View comments
- **read:template:confluence** – View content templates
- **read:label:confluence** – View labels
- **read:watcher:confluence** – View content watchers
- **read:group:confluence** – View groups
- **read:relation:confluence** – View entity relationships
- **read:user:confluence** – View user details
- **read:configuration:confluence** – View Confluence settings
- **read:space:confluence** – View space details
- **read:space.permission:confluence** – View space permissions
- **read:space.property:confluence** – View space properties
- **read:user.property:confluence** – View user properties
- **read:space.setting:confluence** – View space settings
- **read:analytics.content:confluence** – View analytics for content
- **read:content.permission:confluence** – Check content permissions
- **read:content.property:confluence** – View content properties
- **read:content.restriction:confluence** – View content restrictions
- **read:content.metadata:confluence** – View content summaries
- **read:inlinetask:confluence** – View tasks
- **read:task:confluence** – View tasks
- **read:permission:confluence** – View content restrictions and space permissions
- **read:whiteboard:confluence** – View whiteboards
- **read:app-data:confluence** – Read app data

For more information, see [Implementing OAuth 2.0 \(3LO\)](#) and [Determining the scopes required for an operation](#) in Atlassian Developer.

Step 3: Retrieving Confluence (Cloud) client ID and client Secret

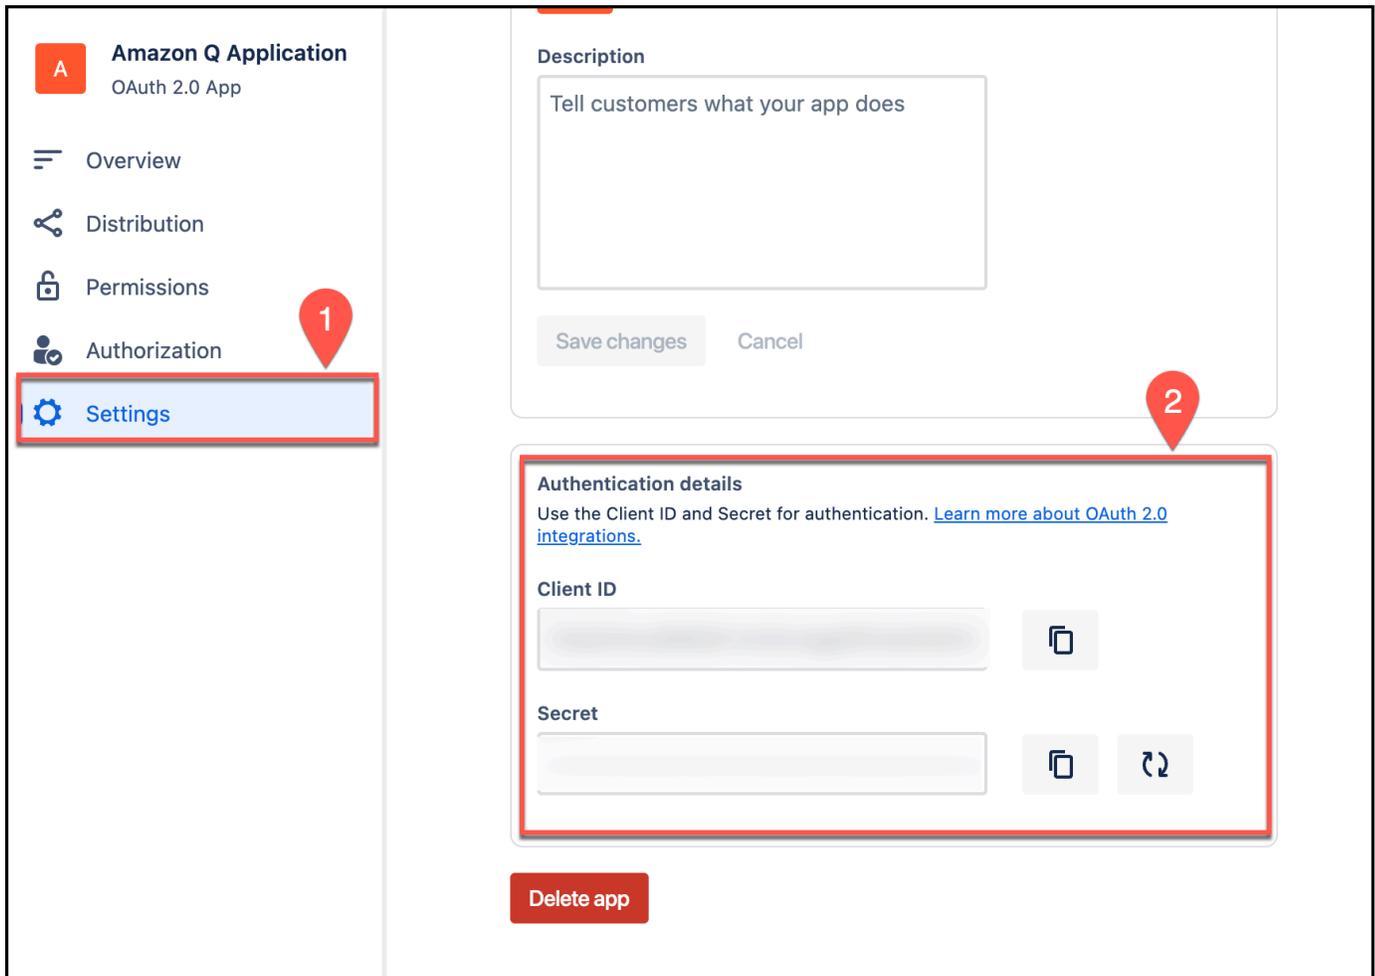
To connect Confluence (Cloud) to Amazon Q using OAuth 2.0 authentication, you need to provide a Confluence (Cloud) client ID and client secret. The following procedure shows you how to retrieve these.

Note

You must create an OAuth 2.0 app before you can retrieve the client ID and client secret. See [Configuring an OAuth 2.0 app integration](#) for more details.

Retrieving Confluence (Cloud) client ID and client secret

- From the left navigation menu, choose **Settings**. Then, scroll down to **Authentication details** section and copy and save the following in a text editor of your choice:
 - Client ID – You will enter this as **App key** in the Amazon Q console.
 - Client Secret – You will enter this as **App secret** in the Amazon Q console.



You will need these to generate your Confluence (Cloud) OAuth 2.0 token and also to connect Amazon Q to Confluence (Cloud).

For more information, see [Implementing OAuth 2.0 \(3LO\)](#) and [Determining the scopes required for an operation](#) in Atlassian Developer.

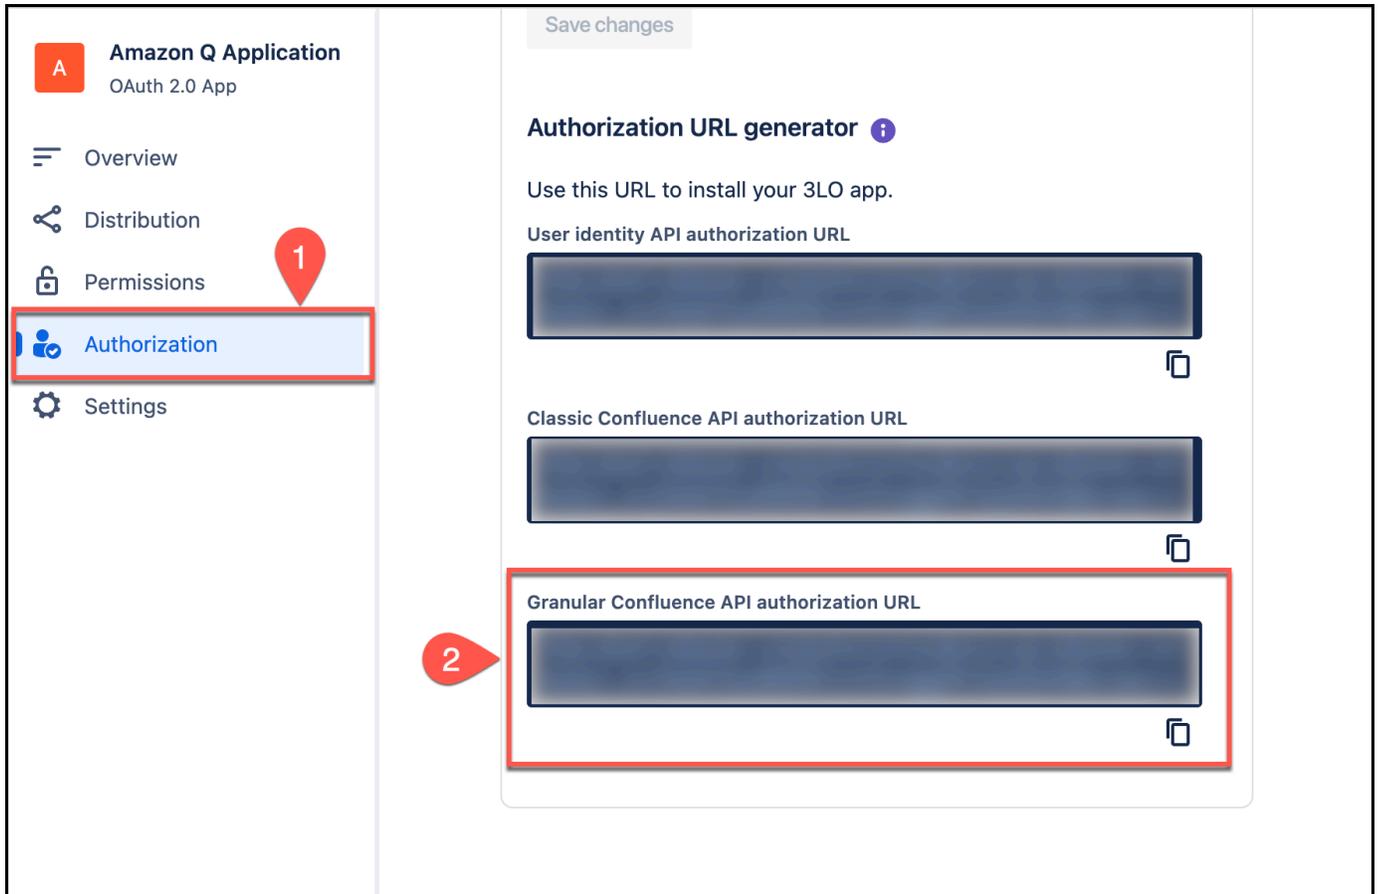
Step 4: Generating an Confluence (Cloud) access token

To connect Confluence (Cloud) to Amazon Q, you need to generate an access token. The following procedure outlines how to generate an access token in Confluence (Cloud).

Generating your Confluence (Cloud) access token

1. Log in to your account from the [Atlassian Developer page](#).
2. Open the OAuth 2.0 app you want to generate a refresh token for.

- From the left navigation menu, choose **Authorization** again. Then, for **OAuth 2.0 (3LO)**, choose **Configure**.
- From the **Authorization** page, from **Authorization URL generator**, from **Granular Confluence API authorization URL**, copy the URL and save it in a text editor of your choice.



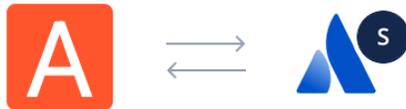
The URL is of the following format:

```
https://auth.atlassian.com/authorize?
audience=api.atlassian.com
&client_id=YOUR_CLIENT_ID
&scope=REQUESTED_SCOPE%20REQUESTED_SCOPE_TWO
&redirect_uri=https://YOUR_APP_CALLBACK_URL
&state=YOUR_USER_BOUND_VALUE
&response_type=code
&prompt=consent
```

- In the saved authorization URL, update the `state=${YOUR_USER_BOUND_VALUE}` parameter value to any text of your choice. For example, `state=sample_text`.

For more information, see [What is the state parameter used for?](#) in Atlassian Support.

6. Open a web browser of your choice. Then, paste the authorization URL you copied into the browser URL. On the page that opens up, make sure everything is correct and then select **Accept**.



[Redacted] is requesting access
to your Atlassian account.

Use app on



 In Confluence, it would like to:

View

- › Analytics for content, App Properties, Content attachments, Audit log records, Blogpost, Comments, Confluence settings, Content metadata, Content Permission, Content Property, Content restrictions, Custom content, Groups, Inline tasks, Labels, Page, Task, Entity relationships, Space permissions, Space properties, Space settings, Space details, Task, Content templates, User properties, User details, Content watchers, Whiteboard

 **Make sure you trust** [Redacted]

This app is in development mode. Development mode apps may pose a risk to your personal data. Only proceed if you know and trust the developer. You can always see and remove access in your Atlassian account.

By accepting this app, you:

- Grant the app access to your data in all places you can access where the app is installed.

This 3rd party vendor has not provided a privacy policy. Atlassian's privacy policy is not applicable to the use of this app.

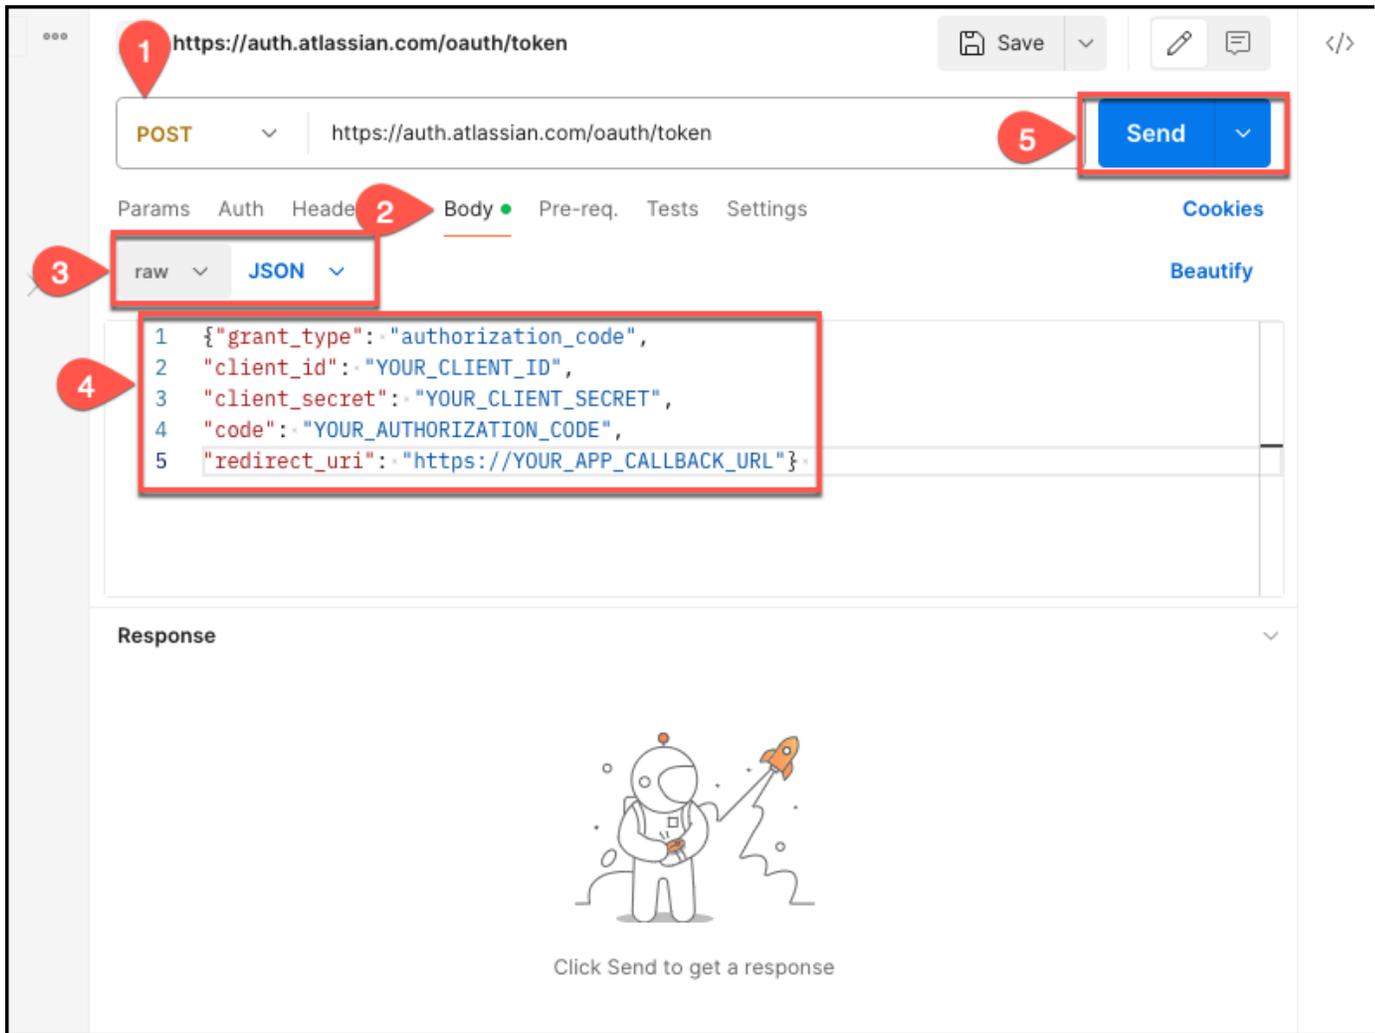
You will be returned to your Confluence (Cloud) home page.

7. Copy the URL of the Confluence (Cloud) home page and save it in a text editor of your choice. The URL contains the authorization code for your application. You will need this code to generate your Confluence (Cloud) access token. The whole section after code= is the authorization code.
8. Navigate to Postman.

If you don't have Postman, you can also choose to use cURL to generate a Confluence (Cloud) access token. Use the following cURL command to do so:

```
curl --location 'https://auth.atlassian.com/oauth/token' \  
--header 'Content-Type: application/json' \  
--data '{"grant_type": "authorization_code",  
"client_id": "YOUR_CLIENT_ID",  
"client_secret": "YOUR_CLIENT_SECRET",  
"code": "AUTHORIZATION_CODE",  
"redirect_uri": "YOUR_CALLBACK_URL"}'
```

9. On the Postman home page, select POST as the method, and then enter the following URL in the **Enter URL or paste text** box: `https://auth.atlassian.com/oauth/token`.
10. Then, select **Body** from the menu, and select **raw JSON**.



11. In the text box, enter the following code extract, replacing the fields with your credential values:

```
{"grant_type": "authorization_code",
"client_id": "YOUR_CLIENT_ID",
"client_secret": "YOUR_CLIENT_SECRET",
"code": "YOUR_AUTHORIZATION_CODE",
"redirect_uri": "https://YOUR_APP_CALLBACK_URL"}
```

12. Then, select **Send**. If everything is configured correctly, Postman will return an access-token. Copy the access token and save it using a text editor of your choice. You will need it to connect Confluence (Cloud) to Amazon Q.

For more information, see [Implementing OAuth 2.0 \(3LO\)](#) in Atlassian Developer.

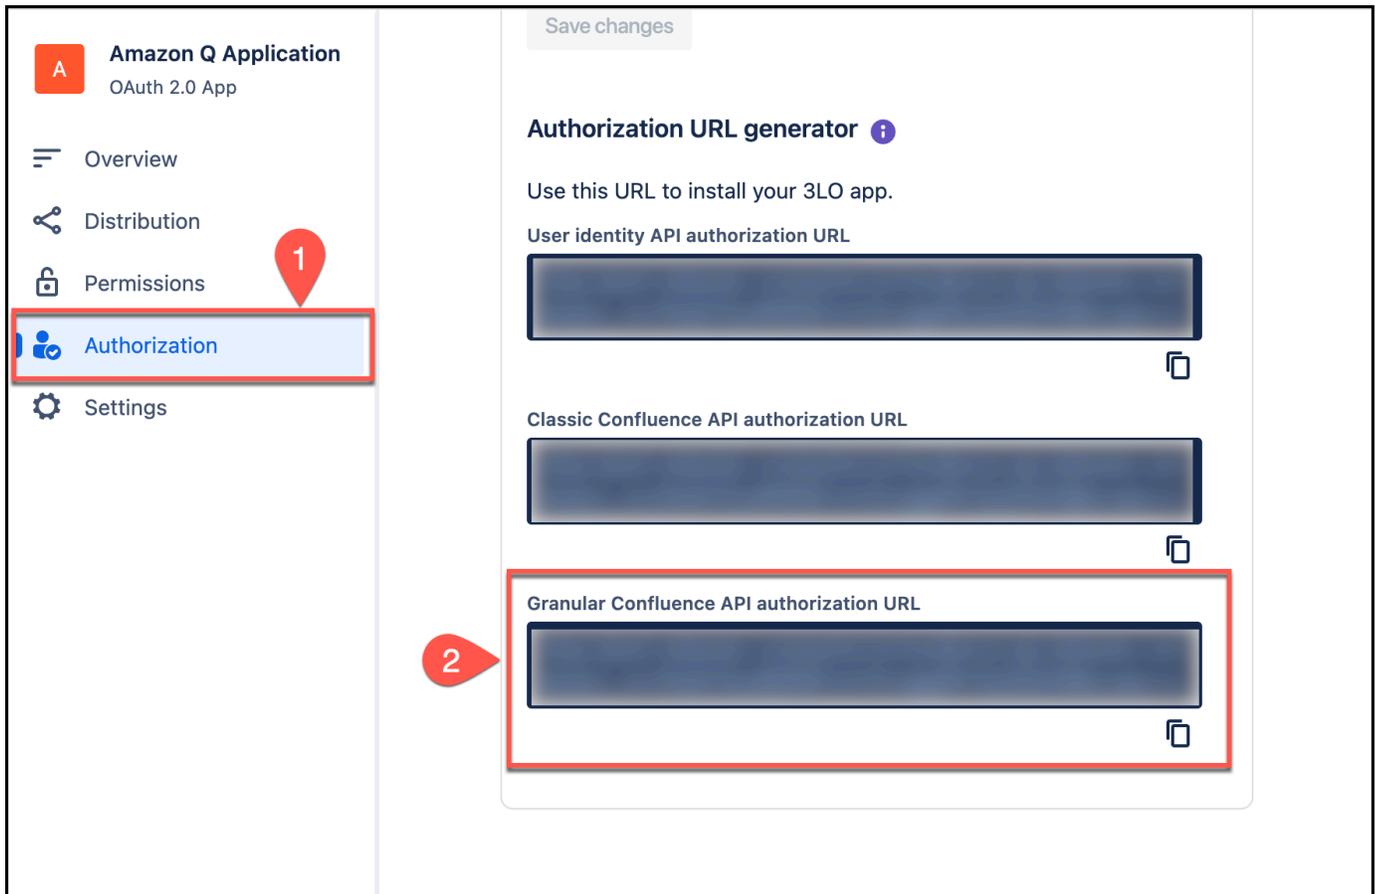
Step 5: Generating a Confluence (Cloud) refresh token

The access token you use to connect Confluence (Cloud) to Amazon Q using OAuth 2.0 authentication expires after 1 hour. When it does, you can either repeat the whole authorization process and generate a new access token. Or, you can choose to generate a refresh token. You can use the refresh token to regenerate a new access token when an existing access token expires.

To do this, you add a `%20offline_access` parameter to the end of the scope value in the authorization URL you used to generate your access token. The following procedure shows you how to generate a refresh token.

Generating an Confluence (Cloud) refresh token

1. Log in to your account from the [Atlassian Developer page](#).
2. Open the OAuth 2.0 app you want to generate a refresh token for.
3. From the left navigation menu, choose **Authorization** again. Then, for **OAuth 2.0 (3LO)**, choose **Configure**.
4. From the **Authorization** page, from **Authorization URL generator**, from **Granular Confluence API authorization URL**, copy the URL and save it in a text editor of your choice.



- In the saved authorization URL, update the `state=${YOUR_USER_BOUND_VALUE}` parameter value to any text of your choice. For example, `state=sample_text`.

For more information, see [What is the state parameter used for?](#) in Atlassian Support.

- Then, add the following text at the end of the scope value in your authorization URL: `%20offline_access` and copy it. For example:

```
https://auth.atlassian.com/authorize?
audience=api.atlassian.com
&client_id=YOUR_CLIENT_ID
&scope=REQUESTED_SCOPE%20REQUESTED_SCOPE_TWO%20offline_access
&redirect_uri=https://YOUR_APP_CALLBACK_URL
&state=YOUR_USER_BOUND_VALUE
&response_type=code
&prompt=consent
```

- Open a web browser of your choice and paste the modified authorization URL you copied into the browser URL. On the page that opens up, make sure everything is correct and then select **Accept**.



[Redacted] is requesting access
to your Atlassian account.

Use app on



 In Confluence, it would like to:

View

- › Analytics for content, App Properties, Content attachments, Audit log records, Blogpost, Comments, Confluence settings, Content metadata, Content Permission, Content Property, Content restrictions, Custom content, Groups, Inline tasks, Labels, Page, Task, Entity relationships, Space permissions, Space properties, Space settings, Space details, Task, Content templates, User properties, User details, Content watchers, Whiteboard

 **Make sure you trust** [Redacted]

This app is in development mode. Development mode apps may pose a risk to your personal data. Only proceed if you know and trust the developer. You can always see and remove access in your Atlassian account.

By accepting this app, you:

- Grant the app access to your data in all places you can access where the app is installed.

This 3rd party vendor has not provided a privacy policy. Atlassian's privacy policy is not applicable to the use of this app.

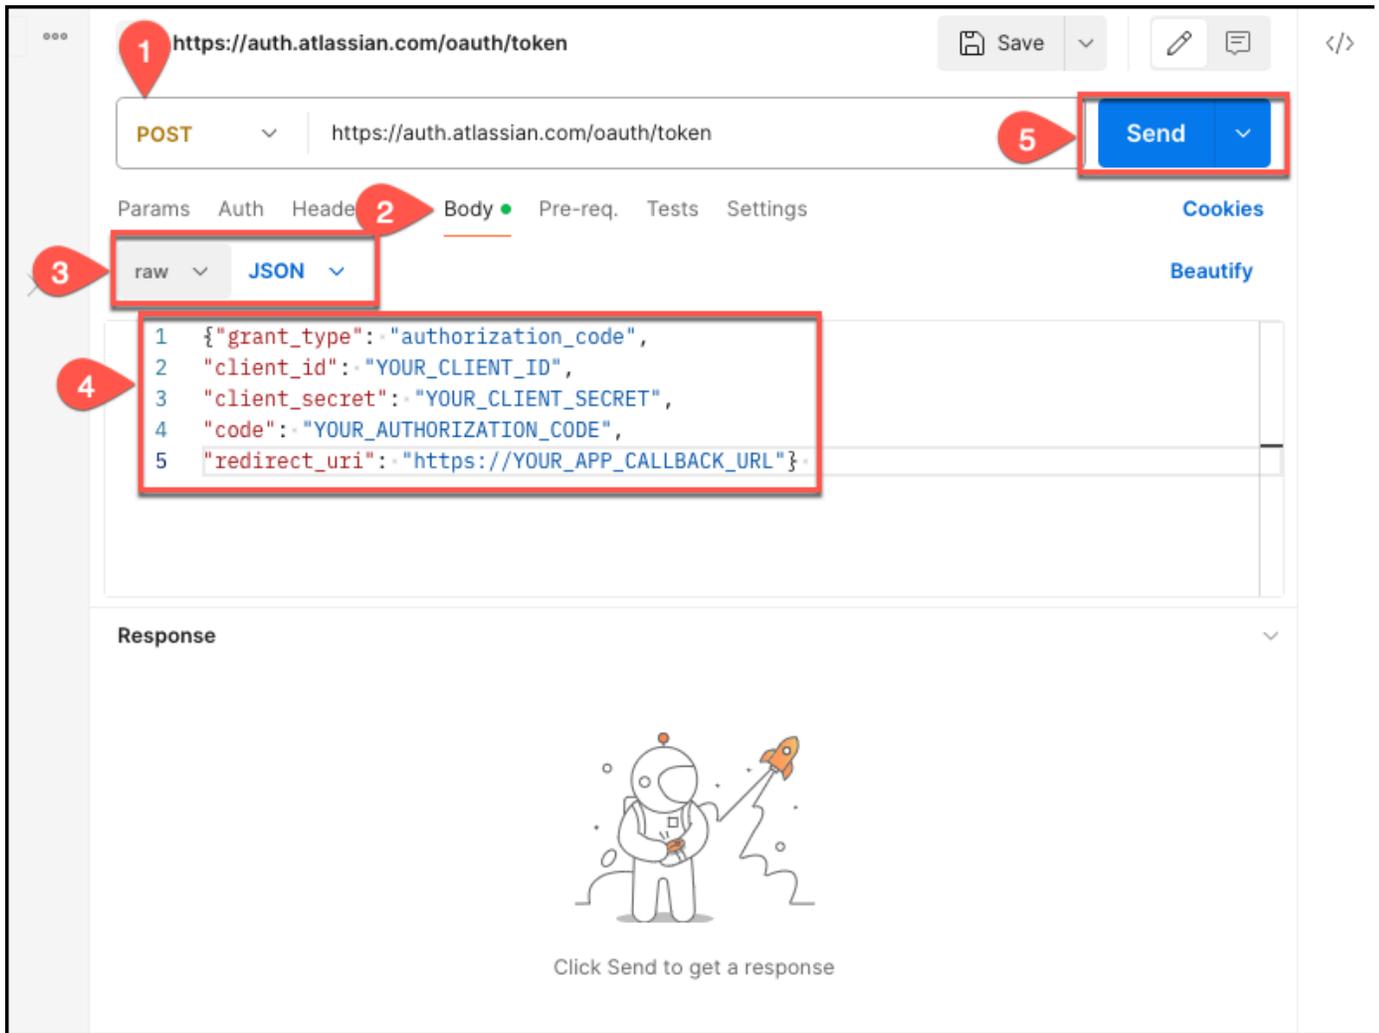
You will be returned to the Confluence (Cloud) console.

8. Copy the URL of the Confluence (Cloud) home page and save it in a text editor of your choice. The URL contains the authorization code for your application. You will need this code to generate your Confluence (Cloud) refresh token. The whole section after code= is the authorization code.
9. Navigate to Postman.

If you don't have Postman, you can also choose to use cURL to generate a Confluence (Cloud) access token. Use the following cURL command to do so:

```
curl --location 'https://auth.atlassian.com/oauth/token' \  
--header 'Content-Type: application/json' \  
--data '{"grant_type": "authorization_code",  
"client_id": "YOUR CLIENT ID",  
"client_secret": "YOUR CLIENT SECRET",  
"code": "AUTHORIZATION CODE",  
"redirect_uri": "YOUR CALLBACK URL"}'
```

10. On the Postman home page, select POST as the method, and then enter the following URL in the **Enter URL or paste text** box: `https://auth.atlassian.com/oauth/token`.
11. Then, select **Body** from the menu, and select **raw JSON**.



12. In the text box, enter the following code extract, replacing the fields with your credential values:

```
{"grant_type": "authorization_code",
"client_id": "YOUR_CLIENT_ID",
"client_secret": "YOUR_CLIENT_SECRET",
"code": "YOUR_AUTHORIZATION_CODE",
"redirect_uri": "https://YOUR_APP_CALLBACK_URL"}
```

13. Then, select **Send**. If everything is configured correctly, Postman will return an refresh-token.

Copy the refresh token and save it using a text editor of your choice. You will need it to connect Confluence (Cloud) to Amazon Q.

For more information, see [Implementing a Refresh Token Flow](#) in Atlassian Developer.

Step 6: Generating a new Confluence (Cloud) access token using a refresh token

You can use the refresh token you generated to create a new access token-refresh token pair when an existing access token expires. The following procedure shows you how to generate a refresh token.

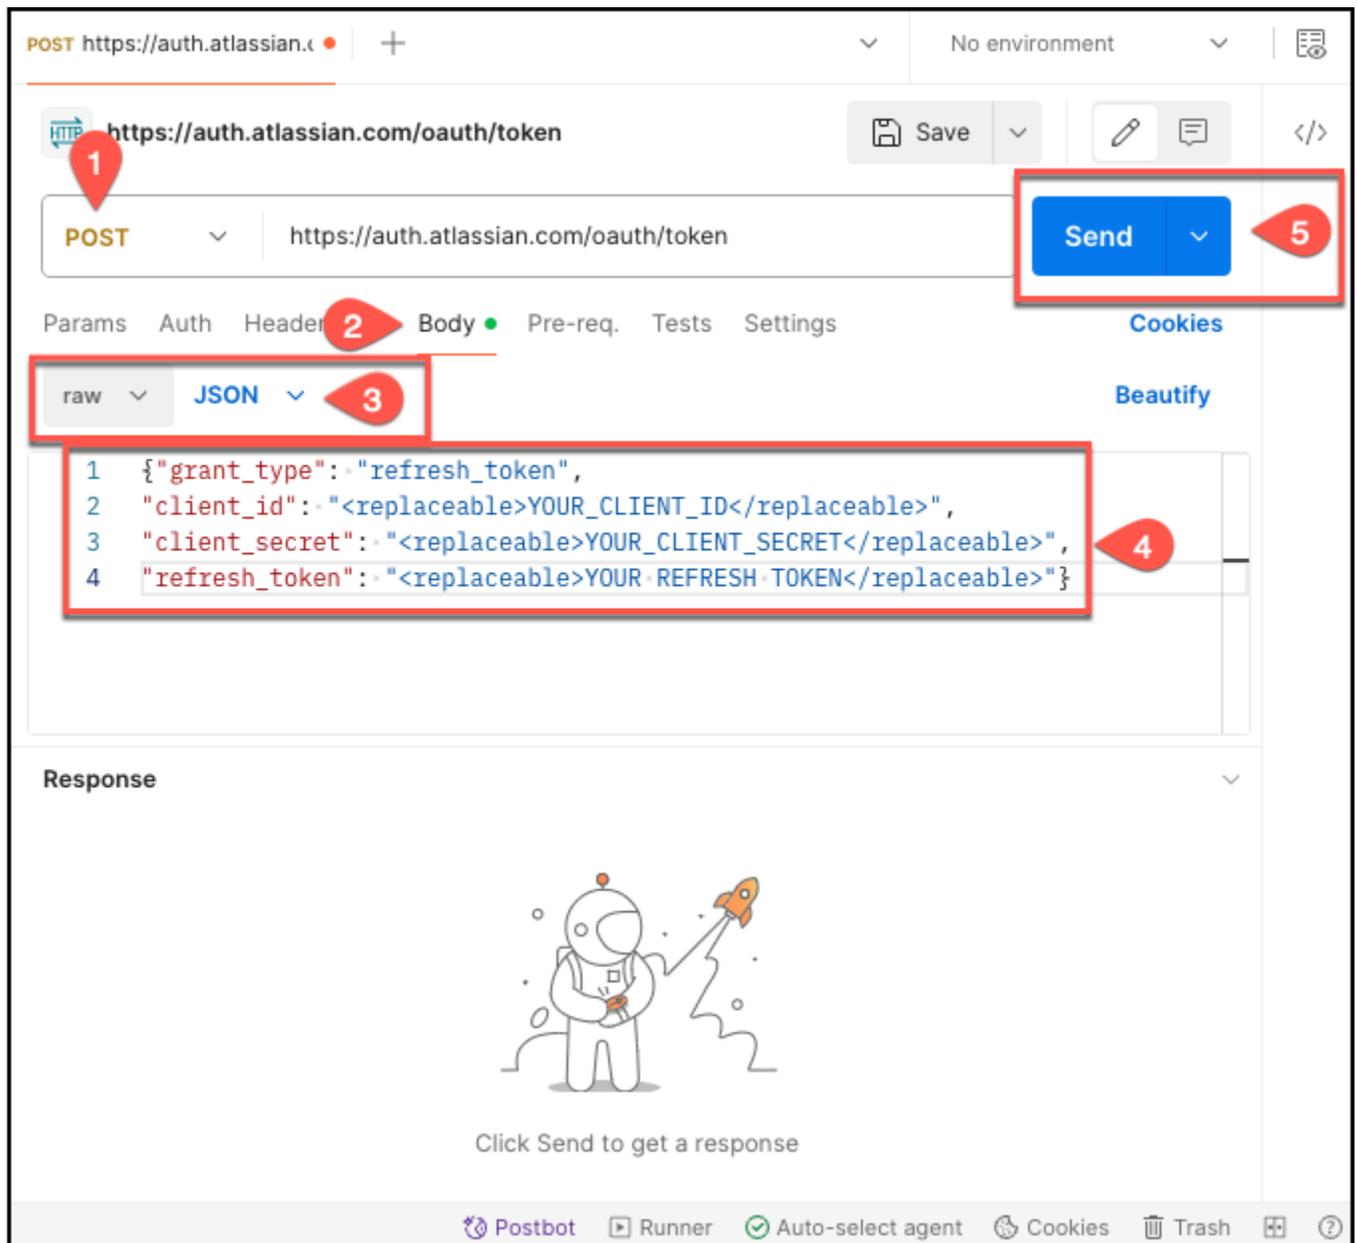
Generating an Confluence (Cloud) access token-refresh token pair

1. Copy the refresh token you generated following the steps in [Step 5: Generating a Confluence \(Cloud\) refresh token](#).
2. Navigate to Postman.

If you don't have Postman, you can also choose to use cURL to generate a new Confluence (Cloud) access token. Use the following cURL command to do so:

```
curl --location 'https://auth.atlassian.com/oauth/token' \  
--header 'Content-Type: application/json' \  
--data '{"grant_type": "refresh_token",  
"client_id": "YOUR_CLIENT_ID",  
"client_secret": "YOUR_CLIENT_SECRET",  
"refresh_token": "YOUR_REFRESH_TOKEN"}'
```

3. On the Postman home page, select POST as the method, and then enter the following URL in the **Enter URL or paste text** box: `https://auth.atlassian.com/oauth/token`.
4. Then, select **Body** from the menu, and select **raw JSON**.



5. In the text box, enter the following code extract, replacing the fields with your credential values:

```
{"grant_type": "refresh_token",  
"client_id": "YOUR_CLIENT_ID",  
"client_secret": "YOUR_CLIENT_SECRET",  
"refresh_token": "YOUR_REFRESH_TOKEN"}
```

6. Then, select **Send**. If everything is configured correctly, Postman will return a new access token-refresh token pair in the following format:

```
{
  "access_token": "string",
  "expires_in": "expiry time of access_token in second",
  "scope": "string",
  "refresh_token": "string"
}
```

For more information, see [Implementing a Refresh Token Flow](#) and [How do I get a new access token, if my access token expires or is revoked?](#) in Atlassian Developer.

How Amazon Q works with Confluence (Cloud) access and refresh tokens

The following are important points to note about using Confluence (Cloud) access and refresh tokens with Amazon Q:

- If a Confluence (Cloud) access token-refresh token pair you use to connect to Amazon Q are expired or invalid, the Amazon Q sync process fails. If this happens, you need to generate and provide a new pair of tokens.
- If your access token is valid but you have an invalid refresh token, Amazon Q will sync data until the access token expires (upto 1 hour). After the access token expires, you won't be able to regenerate an access token-refresh token pair using the expired refresh token. When both access token and refresh token expire, the Amazon Q Confluence (Cloud) data source connector stops syncing.
- If an access token expires during the Confluence (Cloud) connector sync process, the connector internally regenerates a new pair of tokens using the existing refresh token (if the provided refresh token is valid). After regenerating the new pair of tokens, the old pair is invalidated by Confluence (Cloud) and can't be re-used. To sync documents again after the connector auto-regenerates tokens, you must provide a new access token-refresh token pair.
- As a best practice, use the Confluence (Cloud) OAuth app and the generated pair of tokens for the Amazon Q connector only.

Connecting Amazon Q Business to Confluence (Cloud) using the console

The following procedure outlines how to connect Amazon Q Business to Confluence (Cloud) using the AWS Management Console.

Connecting Amazon Q to Confluence (Cloud)

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **Confluence (Cloud)** page, enter the following information:
6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. In **Source**, enter the following information:
 - a. In **Source**, for **Hosting Method**, choose **Confluence Cloud**.
 - b. **Confluence URL** – Enter the Confluence host URLs. The format for the host URL that you enter is *https://example.atlassian.net*.

Important

If you change or update your Confluence (Cloud) data source URL, you also need to update your Secrets Manager secret to ensure a secure connection.

8. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

Note

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

9. For **Authentication** – Choose between **Basic authentication** and **Oauth 2.0 authentication**, based on your use case.

10. **AWS Secrets Manager secret** – Choose an existing secret or create a Secrets Manager secret to store your Confluence authentication credentials. If you choose to create a secret, an AWS Secrets Manager secret window opens. Enter the following information in the window:
 - a. **Secret name** – A name for your secret.
 - b. If using **Basic Authentication** – Enter the **Secret name User name**, and **Password** (Confluence API token) that you generated and downloaded from your Confluence account.

If using **OAuth2.0 Authentication** – Enter the **Secret name**, **App key**, **App secret**, **Access token**, and **Refresh token** that you created in your Confluence account.
 - c. Choose **Save and add secret**.
11. **Configure VPC and security group – optional** – Choose whether you want to use a VPC. If you do, enter the following information:
 - a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
 - b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

12. **Identity crawler** – Choose to activate Amazon Q identity crawler to sync identity information. For more information, see [Identity crawler](#).
13. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

14. In **Sync scope**, choose from the following options :
 - a. In **Sync scope**, for **sync contents**, choose to sync from the following entity types: **Pages**, **Page comments**, **Page attachments**, **Blogs**, **Blog comments**, **Blog attachments**, **Personal spaces Archived spaces**, and **Archived pages**.

Note

Page comments and **Page attachments** can only be selected if you choose to sync **Pages**. **Blog comments** and **Blog attachments** can only be selected if you choose to sync **Blogs**.

Important

You can crawl **Pages** and **Blogs** from one of more specific **Spaces**. If you don't specify a **Space key** regex pattern in **Additional configuration**, all **Pages** and **Blogs** will be crawled by default. If no **Space** is specified in the filter, all spaces will be crawled.

b. In **Additional configuration – optional**, for **Space and regex patterns**, specify whether to include or exclude specific spaces in your index using:

- **Space key** – For example, *my-space-123*.

Note

If you don't specify a **Space key** regex pattern in **Additional configuration**, all **Pages** and **Blogs** will be crawled by default. If no **Space** is specified in the filter, all spaces will be crawled.

- **URL** – For example, *.*MySite/MyDocuments/*.
- **File type** – For example, *.*\pdf, .*\.txt*.
- For **Maximum file size** – Specify the file size limit in MBs that Amazon Q will crawl. Amazon Q will crawl only the files within the size limit you define. The default file size is 50MB. The maximum file size should be greater than 0MB and less than or equal to 50MB.
- For **Entity title regex patterns** – Specify regular expression patterns to include or exclude certain **Blogs**, **Pages**, **Comments**, and **Attachments** by titles.

Note

If you want to crawl a specific page or subpage, you can use page title regex patterns to either include or exclude this page.

15. For **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.
 - **Full sync** – Sync all content regardless of the previous sync status.
 - **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.
16. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
17. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
18. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:
 - a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
 - b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

Note

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

19. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

20. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

Note

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to Confluence (Cloud) using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the `configuration` parameter to provide a JSON schema with all other configuration information specific to your data source connector.

For an example of the API request, see [CreateDataSource](#) in the Amazon Q API Reference.

Confluence JSON schema

The following is the Confluence JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
      "type": "object",
      "properties": {
        "repositoryEndpointMetadata": {
          "type": "object",
          "properties": {
            "hostUrl": {
              "type": "string",
              "pattern": "https:.*"
            },
            "type": {
              "type": "string",
              "enum": [
```

```

        "SAAS",
        "ON_PREM"
    ]
},
"authType": {
    "type": "string",
    "enum": [
        "Basic",
        "OAuth2",
        "Personal-token"
    ]
},
"required": [
    "hostUrl",
    "type",
    "authType"
]
},
"required": [
    "repositoryEndpointMetadata"
]
},
"repositoryConfigurations": {
    "type": "object",
    "properties": {
        "space": {
            "type": "object",
            "properties": {
                "fieldMappings": {
                    "type": "array",
                    "items": [
                        {
                            "type": "object",
                            "properties": {
                                "indexFieldName": {
                                    "type": "string"
                                },
                                "indexFieldType": {
                                    "type": "string",
                                    "enum": [
                                        "STRING",
                                        "STRING_LIST",

```



```

        "DATE"
      ]
    },
    "dataSourceFieldName": {
      "type": "string"
    },
    "dateFieldFormat": {
      "type": "string",
      "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
    }
  },
  "required": [
    "indexFieldName",
    "indexFieldType",
    "dataSourceFieldName"
  ]
}
]
}
},
"required": [
  "fieldMappings"
]
},
"page": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": [
        {
          "type": "object",
          "properties": {
            "indexFieldName": {
              "type": "string"
            },
            "indexFieldType": {
              "type": "string",
              "enum": [
                "STRING",
                "STRING_LIST",
                "DATE",
                "LONG"
              ]
            }
          }
        }
      ]
    }
  }
}

```

```

    },
    "dataSourceFieldName": {
      "type": "string"
    },
    "dateFieldFormat": {
      "type": "string",
      "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
    }
  },
  "required": [
    "indexFieldName",
    "indexFieldType",
    "dataSourceFieldName"
  ]
}
]
}
},
"required": [
  "fieldMappings"
]
},
"blog": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": [
        {
          "type": "object",
          "properties": {
            "indexFieldName": {
              "type": "string"
            },
            "indexFieldType": {
              "type": "string",
              "enum": [
                "STRING",
                "STRING_LIST",
                "DATE",
                "LONG"
              ]
            }
          }
        }
      ]
    },
    "dataSourceFieldName": {

```

```

        "type": "string"
      },
      "dateFieldFormat": {
        "type": "string",
        "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
      }
    },
    "required": [
      "indexFieldName",
      "indexFieldType",
      "dataSourceFieldName"
    ]
  }
]
}
},
"required": [
  "fieldMappings"
]
},
"comment": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": [
        {
          "type": "object",
          "properties": {
            "indexFieldName": {
              "type": "string"
            },
            "indexFieldType": {
              "type": "string",
              "enum": [
                "STRING",
                "STRING_LIST",
                "DATE",
                "LONG"
              ]
            },
            "dataSourceFieldName": {
              "type": "string"
            }
          }
        }
      ]
    }
  }
}

```

```

        "dateFieldFormat": {
            "type": "string",
            "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
        }
    },
    "required": [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
    ]
}
]
}
},
"required": [
    "fieldMappings"
]
},
"attachment": {
    "type": "object",
    "properties": {
        "fieldMappings": {
            "type": "array",
            "items": [
                {
                    "type": "object",
                    "properties": {
                        "indexFieldName": {
                            "type": "string"
                        },
                        "indexFieldType": {
                            "type": "string",
                            "enum": [
                                "STRING",
                                "STRING_LIST",
                                "DATE",
                                "LONG"
                            ]
                        },
                        "dataSourceFieldName": {
                            "type": "string"
                        },
                        "dateFieldFormat": {
                            "type": "string",

```

```

        "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
      }
    },
    "required": [
      "indexFieldName",
      "indexFieldType",
      "dataSourceFieldName"
    ]
  }
]
}
},
"required": [
  "fieldMappings"
]
}
},
"additionalProperties": {
  "type": "object",
  "properties": {
    "isCrawlAcl": {
      "type": "boolean"
    },
    "fieldForUserId": {
      "type": "string"
    },
    "inclusionSpaceKeyFilter": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    "exclusionSpaceKeyFilter": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    "pageTitleRegEX": {
      "type": "array",
      "items": {
        "type": "string"
      }
    }
  }
}

```

```
  },
  "blogTitleRegEX": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "commentTitleRegEX": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "attachmentTitleRegEX": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "isCrawlPersonalSpace": {
    "type": "boolean"
  },
  "isCrawlArchivedSpace": {
    "type": "boolean"
  },
  "isCrawlArchivedPage": {
    "type": "boolean"
  },
  "isCrawlPage": {
    "type": "boolean"
  },
  "isCrawlBlog": {
    "type": "boolean"
  },
  "isCrawlPageComment": {
    "type": "boolean"
  },
  "isCrawlPageAttachment": {
    "type": "boolean"
  },
  "isCrawlBlogComment": {
    "type": "boolean"
  },
  "isCrawlBlogAttachment": {
```

```
    "type": "boolean"
  },
  "maxFileSizeInMegabytes": {
    "type": "string"
  },
  "inclusionFileTypePatterns": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "exclusionFileTypePatterns": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "inclusionUrlPatterns": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "exclusionUrlPatterns": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "proxyHost": {
    "type": "string"
  },
  "proxyPort": {
    "type": "string"
  }
},
"required": [],
"type": {
  "type": "string",
  "enum": [
    "CONFLUENCEV2",
    "CONFLUENCE"
  ]
}
```

```

    },
    "enableIdentityCrawler": {
      "type": "boolean"
    },
    },
    "syncMode": {
      "type": "string",
      "enum": [
        "FULL_CRAWL",
        "FORCED_FULL_CRAWL"
      ]
    },
    },
    "secretArn": {
      "type": "string",
      "minLength": 20,
      "maxLength": 2048
    }
  },
  "version": {
    "type": "string",
    "anyOf": [
      {
        "pattern": "1.0.0"
      }
    ]
  },
  },
  "required": [
    "connectionConfiguration",
    "repositoryConfigurations",
    "syncMode",
    "additionalProperties",
    "secretArn",
    "type"
  ]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.

Configuration	Description
repositoryEndpointMetadata	The endpoint information for the data source.
hostUrl	<p>The URL for your Confluence instance. For example, <i>https://example.atlassian.net</i> .</p> <div style="border: 1px solid #f08080; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p>⚠ Important</p> <p>If you change or update your Confluence (Cloud) data source URL, you also need to update your Secrets Manager secret to ensure a secure connection.</p> </div>
type	The hosting method for your Confluence instance, whether SAAS or ON_PREM.
authType	The authentication method for your Confluence instance, whether Basic or OAuth2.
repositoryConfigurations	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings.
<ul style="list-style-type: none"> • space • page • blog • comment • attachment 	A list of objects that map the attributes or field names of your Confluence spaces, pages, blogs, comments, and attachments to Amazon Q index field names.
additionalProperties	Additional configuration options for your content in your data source.
isCrawlAcl	Specify true to crawl access control information from documents.

Configuration	Description
fieldForUserId	Specify field to use for UserId for ACL crawling.
proxyHost	The host where the web proxy is required. The host name should be without protocol (http:// or https://).
proxyPort	Port used by the host URL transport protocol. The port number should be a numeric value between 0 and 65535.
maxFileSizeInMegabytes	Specify the file size limit in MBs that Amazon Q will crawl. Amazon Q will crawl only the files within the size limit you define. The default file size is 50MB. The maximum file size should be greater than 0MB and less than or equal to 50MB.
<ul style="list-style-type: none"> • inclusionSpaceKeyFilter • exclusionSpaceKeyFilter • pageTitleRegEX • blogTitleRegEX • commentTitleRegEX • attachmentTitleRegEX • inclusionFileTypePatterns • exclusionFileTypePatterns • inclusionUrlPatterns • exclusionUrlPatterns 	A list of regular expression patterns to include and/or exclude certain files in your Confluence data source. Files that match the patterns are included in the index. Files that don't match the patterns are excluded from the index. If a file matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence and the file isn't included in the index.

Configuration	Description
<ul style="list-style-type: none"> • <code>isCrawlPersonalSpace</code> • <code>isCrawlArchivedSpace</code> • <code>isCrawlArchivedPage</code> • <code>isCrawlPage</code> • <code>isCrawlBlog</code> • <code>isCrawlPageComment</code> • <code>isCrawlPageAttachment</code> • <code>isCrawlBlogComment</code> • <code>isCrawlBlogAttachment</code> 	<p><code>true</code> to index files in your Confluence personal spaces, pages, blogs, page comments, page attachments, blog comments, and blog attachments.</p>
<p><code>type</code></p>	<p>The type of data source. Specify <code>CONFLUENCE_V2</code> as your data source type.</p>
<p><code>enableIdentityCrawler</code></p>	<p><code>true</code> to activate identity crawler. Identity crawler is activated by default. See Identity crawler for more information.</p>
<p><code>syncMode</code></p>	<p>Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose between the following options:</p> <ul style="list-style-type: none"> • Use <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index • Use <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index

Configuration	Description
secretARN	<p>The Amazon Resource Name (ARN) of a Secrets Manager secret that contains the key-value pairs required to connect to your Confluence instance.</p> <p>If you use basic authentication, the secret must contain a JSON structure with the following keys:</p> <pre data-bbox="829 615 1507 1014">{ "hostUrl": " <i>Confluence Cloud</i> <i>host URL</i>", "username": " <i>Confluence account</i> <i>user name</i>", "password": " <i>Confluence API</i> <i>token</i>" }</pre> <p>If you use OAuth 2.0 authentication, the secret must contain a JSON structure with the following keys:</p> <pre data-bbox="829 1220 1507 1850">{ "hostUrl": " <i>Confluence Cloud</i> <i>host URL</i>", "confluenceAppKey": " <i>client ID</i> <i>for your Confluence account</i> ", "confluenceAppSecret": " <i>client</i> <i>secret from your Confluence</i> <i>account</i>", "confluenceAccessToken": " <i>access</i> <i>token created in Confluence</i> ", "confluenceRefreshToken": "<i>refresh token created in Confluenc</i> <i>e</i> ", }</pre>

Configuration	Description
version	The version of this template that's currently supported.

How Amazon Q Business connector crawls Confluence (Cloud) ACLs

When you connect an Confluence (Cloud) data source to Amazon Q Business, Amazon Q crawls ACL information attached to a document (user and group information) from your Confluence (Cloud) instance. If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

You configure user and group access to spaces using the space permissions page. For pages and blogs, you use the restrictions page. For more information about space permissions, see [Space Permissions Overview](#) on the Confluence Support website. For more information about page and blog restrictions, see [Page Restrictions](#) on the Confluence Support website.

The group and user IDs are mapped as follows:

- `_group_ids` – Group names are present on spaces, pages, and blogs where there are restrictions. They're mapped from the name of the group in Confluence. Group names are always lower case.
- `_user_id` – User names are present on the space, page, or blog where there are restrictions. They're mapped depending on the type of Confluence instance that you are using.
- For Confluence Cloud – The `_user_id` is the account ID of the user.

Important

For user context filtering to work correctly for your Confluence connector, you need to make sure that the visibility of a user granted access to a Confluence page is set to **Anyone**. For more information, see [Set your email visibility](#) in Atlassian Developer Documentation.

For more information, see:

- [Authorization](#)
- [Identity crawler](#)

- [Understanding User Store](#)

Amazon Q Business Confluence (Cloud) data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q Business enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a attribute mapping already available, or if you want to map additional document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q Confluence connector supports the following entities and the associated reserved and custom attributes.

Supported entities and field mappings

- [Space](#)
- [Page](#)

- [Blog](#)
- [Comment](#)
- [Attachment](#)

Space

Confluence field name	Index field name	Description	Data type
spaceName	cf_sp_document_title	Custom	String
itemType	_category	Default	String
url	_source_uri	Default	String
spaceKey	cf_space_key	Custom	String
description	cf_description	Custom	String
spaceType	cf_type	Custom	String

Page

Confluence field name	Index field name	Description	Data type
title	_cf_page_document_title	Custom	String
authors	_authors	Default	String list
createdDate	_created_at	Default	Date
modifiedDate	_last_updated_at	Default	Date
labels	cf_labels	Custom	String list
version	cf_version	Custom	Long (numeric)

Confluence field name	Index field name	Description	Data type
itemType	_category	Default	String
spaceKey	cf_space_key	Custom	String
spaceName	cf_space_name	Custom	String
url	_source_uri	Default	String
status	cf_status	Custom	String
parentId	cf_parent_id	Custom	String

Blog

Confluence field name	Index field name	Description	Data type
title	cf_bg_document_title	Custom	String
author	_authors	Default	String list
publishedDate	_created_at	Default	Date
labels	_source_uri	Default	String
version	cf_version	Custom	Long (numeric)
itemType	_category	Custom	String
spaceKey	cf_space_key	Custom	String
modifiedDate	_last_updated_at	Default	Date
spaceName	cf_space_name	Custom	String
status	cf_status	Custom	String

Confluence field name	Index field name	Description	Data type
url	_source_uri	Default	String
parentId	cf_parent_id	Custom	String

Comment

Confluence field name	Index field name	Description	Data type
title	cf_cmt_document_title	Custom	String
author	_authors	Default	String list
createdDate	_created_at	Default	Date
version	cf_version	Custom	Long (numeric)
itemType	_category	Default	String
spaceKey	cf_space_key	Custom	String
spaceName	cf_space_name	Custom	String
contentType	cf_content_type	Custom	String
url	_source_uri	Default	String
parentId	cf_parent_id	Custom	String
status	cf_status	Custom	String

Attachment

Confluence field name	Index field name	Description	Data type
fileName	cf_attachment_document_title	Custom	String
author	_authors	Default	String list
createdDate	_created_at	Default	Date
labels	cf_labels	Custom	String list
version	cf_version	Custom	Long (numeric)
itemType	_category	Default	String
spaceKey	cf_space_key	Custom	String
contentType	cf_content_type	Custom	String
modifiedDate	_last_updated_at	Default	Date
fileSize	cf_file_size	Custom	Long (numeric)
fileType	cf_attachment_file_type	Custom	String
spaceName	cf_space_name	Custom	String
documentId	_document_id	Default	String list
url	_source_uri	Default	String
parentId	cf_parent_id	Custom	String
attachmentComment	cf_attachment_comment	Custom	String
status	cf_status	Custom	String

IAM role for Amazon Q Confluence (Cloud) connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the `BatchPutDocument` and `BatchDeleteDocument` operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToGetSecret",
      "Effect": "Allow",
      "Action": [
        "secretsmanager:GetSecretValue"
      ],
      "Resource": [
        "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
      ]
    },
    {
      "Sid": "AllowsAmazonQToDecryptSecret",
      "Effect": "Allow",
      "Action": [
        "kms:Decrypt"
      ],
      "Resource": [
```

```

    "arn:aws:kms:{{region}}:{{account_id}}:key/[{{key_id}}]"
  ],
  "Condition": {
    "StringLike": {
      "kms:ViaService": [
        "secretsmanager.*.amazonaws.com"
      ]
    }
  }
},
{
  "Sid": "AllowsAmazonQToIngestDocuments",
  "Effect": "Allow",
  "Action": [
    "qbusiness:BatchPutDocument",
    "qbusiness:BatchDeleteDocument"
  ],
  "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
},
{
  "Sid": "AllowsAmazonQToIngestPrincipalMapping",
  "Effect": "Allow",
  "Action": [
    "qbusiness:PutGroup",
    "qbusiness:CreateUser",
    "qbusiness>DeleteGroup",
    "qbusiness:UpdateUser",
    "qbusiness:ListGroups"
  ],
  "Resource": [
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}",
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}",
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}/data-source/*"
  ]
},
{
  "Sid": "AllowsAmazonQToCreateAndDeleteNI",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateNetworkInterface",
    "ec2>DeleteNetworkInterface"
  ]
}

```

```

    ],
    "Resource": [
      "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
      "arn:aws:ec2:{{region}}:{{account_id}}:security-group/[{{security_group}}]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2:DeleteNetworkInterface"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:RequestTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
      },
      "ForAllValues:StringEquals": {
        "aws:TagKeys": [
          "AMAZON_Q"
        ]
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateTags",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateTags"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringEquals": {
        "ec2:CreateAction": "CreateNetworkInterface"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterfacePermission"
    ],
  },

```

```

"Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
"Condition": {
  "StringLike": {
    "aws:ResourceTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
  }
},
{
  "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
  "Effect": "Allow",
  "Action": [
    "ec2:DescribeNetworkInterfaces",
    "ec2:DescribeAvailabilityZones",
    "ec2:DescribeNetworkInterfaceAttribute",
    "ec2:DescribeVpcs",
    "ec2:DescribeRegions",
    "ec2:DescribeNetworkInterfacePermissions",
    "ec2:DescribeSubnets"
  ],
  "Resource": "*"
}
]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "{{source_account}}"
        },
        "ArnEquals": {
          "aws:SourceArn": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/{{application_id}}"
        }
      }
    }
  ]
}

```

```

    }
  }
}
]
}

```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Troubleshooting your Amazon Q Business Confluence (Cloud) connector

The following table provides information about error codes you may see for the Confluence (Cloud) connector and suggested troubleshooting actions.

Error code	Error message	Suggested resolution
CNF-5500	Null/empty username.	Provide username.
CNF-5501	Error validating credentials due to Invalid username or password.	Provide valid username/password
CNF-5502	Null/empty confluence AppKey.	Provide confluence AppKey.
CNF-5503	Null/empty confluence Secret.	Provide confluence Secret.
CNF-5504	Null/empty Client Access Token.	Provide Client Access Token.
CNF-5505	Null/empty Client Refresh Token	Provide Client Refresh Token
CNF-5506	Incorrect auth type.	Auth type should be Basic or OAuth2 or Personal-token.
CNF-5507	Null/empty auth type.	Auth Type should not be null or empty value.

Error code	Error message	Suggested resolution
CNF-5508	Empty/null host URL.	Host Url should not be null or empty value.
CNF-5509	Null/empty crawl type.	Crawl Type should not be null or empty value.
CNF-5510	Null/empty Repository Configurations.	Repository Configurations should not be null or empty value.
CNF-5511	Incorrect type.	type should be SAAS or ON_PREM.
CNF-5512	Invalid inclusion file type patterns.	Provide the correct inclusion patterns.
CNF-5513	Invalid exclusion file type patterns.	Provide the correct exclusion patterns.
CNF-5514	Invalid regex patterns.	Provide the correct regex patterns.
CNF-5515	Error validating credentials due to invalid username or password.	Provide valid user name and password.
CNF-5516	Error validating credentials due to invalid client id or client secret.	Provide valid client id and client secret.
CNF-5517	Error validating crawl type.	Provide valid crawl type.
CNF-5518	Invalid URI.	Provide valid URI.
CNF-5519	Null/empty DataSourceFieldName in Space Entity.	Provide value for DataSourceFieldName in Space Entity.
CNF-5520	Null/empty IndexFieldName in Blog Entity.	Provide value for IndexFieldName in Blog Entity.

Error code	Error message	Suggested resolution
CNF-5521	Null/empty IndexField dType in Space Entity.	Provide value for IndexFieldType in Space Entity.
CNF-5522	Null/empty password.	Provide password.
CNF-5523	Incorrect auth type.	Auth type should be Basic or OAuth2.
CNF-5524	Null/empty DataSource eFieldName in Page Entity.	Provide value for DataSourceFieldNam e in Page Entity.
CNF-5525	Null/empty DataSource eFieldName in Blog Entity	Please provide value for DataSourc eFieldName in Blog Entity
CNF-5526	Null/empty DataSource eFieldName in Comment Entity.	Provide value for DataSourceFieldNam e in Comment Entity.
CNF-5527	Null/empty DataSource eFieldName in Attachment Entity.	Provide value for DataSourceFieldNam e in Attachment Entity.
CNF-5528	Null/empty IndexFiel dName.	IndexFieldName field can't be null or empty value.
CNF-5529	Null/empty IndexFiel dName in Space Entity.	Provide value for IndexFieldName in Space Entity.
CNF-5530	Null/empty IndexFiel dName in Page Entity	Please provide value for IndexFiel dName in Page Entity
CNF-5531	Invalid isCrawlPe rsonalSpace value.	isCrawlPersonalSpace should be a boolean value true or false.
CNF-5532	Invalid isCrawlAr chivedSpace value.	isCrawlArchivedSpace should be a boolean value true or false.

Error code	Error message	Suggested resolution
CNF-5533	Invalid isCrawlArchivedPage value.	isCrawlArchivedPage should be a boolean value true or false.
CNF-5534	Invalid isCrawlPage value.	isCrawlPage should be a boolean value true or false.
CNF-5535	Invalid isCrawlBlogComment value.	isCrawlBlogComment should be a boolean value true or false.
CNF-5536	Invalid isCrawlBlogComment value.	isCrawlBlogComment should be a boolean value true or false.
CNF-5537	Invalid isCrawlBlogAttachment value.	isCrawlBlogAttachment should be a boolean value true or false.
CNF-5538	Error validating on protocol.	Provide valid protocol.
CNF-5539	Null/empty IndexFieldName in Comment Entity.	Provide value for IndexFieldName in Comment Entity.
CNF-5540	Null/empty Personal Access Token.	Provide Personal Access Token.
CNF-5541	Invalid OAuth value.	Give a valid OAuth URL.
CNF-5542	Invalid Space value.	Give a valid Space URL.
CNF-5543	Archived Space Exception .	Check Archived Space.
CNF-5544	JSON Exception for Space.	Check Space.
CNF-5545	JSON Exception for Comment.	Check Comment.

Error code	Error message	Suggested resolution
CNF-5546	JSON Exception for Comment.	Check Comment.
CNF-5547	JSON Exception for Comment.	Check Comment.
CNF-5548	JSON Exception for Attachment.	Check Attachment.
CNF-5549	JSON Exception for Blog.	Check Blog.
CNF-5550	JSON Exception for Page.	Check Page.
CNF-5551	JSON Exception for Label.	Check Label.
CNF-5552	JSON Exception for ACL.	Check ACL.
CNF-5553	JSON Exception for Groups.	Check Groups.
CNF-5554	JSON Exception for Group Members.	Check Group Members.
CNF-5555	JSON Exception for Space Group.	Check Space Group.
CNF-5556	Exception in CommentItem.	Check the CommentItem class.
CNF-5557	Invalid isCrawlPageComment value.	isCrawlPageComment should be a boolean value true or false.
CNF-5558	Invalid isCrawlPageAttachment value.	isCrawlPageAttachment should be a boolean value true or false.
CNF-5559	Null/empty Repository Configurations.	Repository Configurations should not be null or empty value.

Error code	Error message	Suggested resolution
CNF-5560	Null/empty IndexFieldName in Attachment.	Please provide value for IndexFieldname in Attachment Entity.
CNF-5561	Invalid proxy url.	Proxy url should not contain http: or https.
CNF-5562	Null/Empty proxy port.	Provide a valid proxy port.
CNF-5563	Invalid Host URL.	Provide valid Host URL.
CNF-5564	Invalid proxy port value.	Provide a valid proxy port.
CNF-5565	Confluence server not reachable.	Provide a valid proxy and server details.
CNF-5566	Null/empty IndexFieldtype in Page Entity.	Provide value for IndexFieldType in Page Entity.
CNF-5567	Null/empty IndexFieldtype in Blog Entity.	Provide value for IndexFieldType in Blog Entity.
CNF-5568	Null/empty IndexFieldtype in Comment Entity.	Provide value for IndexFieldType in Comment Entity.
CNF-5569	Null/empty IndexFieldtype in Attachment.	Provide value for IndexFieldType in Attachment. Entity
CNF-5570	JSON Exception for Content Ancestors.	Check your Ancestors.
CNF-5571	Invalid Host URL Pattern.	Provide valid Host URL Pattern.
CNF-5572	Error validating credentials due to Invalid access or refresh token.	Invalid AccessToken/RefreshToken.

Connecting Confluence (Server/Data Center) to Amazon Q Business

Atlassian Confluence is a collaborative work-management tool designed for sharing, storing, and working on project planning, software development, and product management. You can connect Confluence (Server/Data Center) instance to Amazon Q Business—using either the AWS Management Console or the [CreateDataSource](#) API—and create an Amazon Q web experience.

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).
- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [Overview](#)
- [Prerequisites](#)
- [Using the console](#)
- [Connecting Amazon Q Business to Confluence \(Server/Data Center\) using APIs](#)
- [How Amazon Q Business connector crawls Confluence \(Server/Data Center\) ACLs](#)
- [Amazon Q Business Confluence \(Server/Data Center\) data source connector field mappings](#)
- [IAM role for Amazon Q Confluence \(Server/Data Center\) connector](#)
- [Troubleshooting your Amazon Q Business Confluence \(Server/Data Center\) connector](#)

Overview

The following table gives an overview of the Amazon Q Business Confluence (Server/Data Center) connector and its supported features.

Category	Feature	Support
Security	Authentication type	Basic, OAuth 2.0, Personal Access Token
	Authentication credentials	For Basic authentication

Category	Feature	Support
		<ul style="list-style-type: none"> • Confluence Server/Data Center URL • Confluence API token <p>For OAuth 2.0 authentication</p> <ul style="list-style-type: none"> • App key • App secret • Access token • Refresh token <div data-bbox="860 703 1510 1018" style="border: 1px solid #add8e6; border-radius: 10px; padding: 10px; margin: 10px 0;"> <p>Note</p> <p>Access and refresh tokens expire in 1 hour. For information on regenerating tokens, see Atlassian Developer Documentation.</p> </div> <p>Personal Access Token</p> <ul style="list-style-type: none"> • Personal Access Token
	<p>Access Control List (ACL) crawling</p>	<p>Yes. For more information, see ACL crawling.</p>
	<p>Identity crawling</p>	<p>Yes</p>
<p>Crawl features</p>	<p>Custom metadata</p>	<p>Yes</p>

Category	Feature	Support
	Entities	Yes. The following entities are supported: <ul style="list-style-type: none"> • Space • Page • Blog post • Comment • Attachment
	Field mappings	Yes. For more information, see Field mappings .
	Filters	Yes. The following filters are supported: <ul style="list-style-type: none"> • Inclusion exclusion filters for Space key and Space URL • Inclusion exclusion filters on File Type for Attachment entity • Supports regex filters for entities • Supports inclusion and exclusion filters for File size
	Sync mode	Supports full and incremental (new, modified, and deleted) sync.
	File types	Supports all files supported by Amazon Q.

Prerequisites

Before you begin, make sure that you have completed the following prerequisites.

In Confluence Server/Data Center, make sure you have:

- Copied your Confluence instance URL. For example: *https://example.confluence.com*. You need your Confluence instance URL to connect to Amazon Q.
- Configured basic authentication credentials containing a username (email ID used to log into Confluence) and password (Confluence Server/Data Center password) to allow Amazon Q to

connect to your Confluence Server/Data Center instance. For information about how to create a Confluence API token, see [Manage API tokens for your Atlassian account](#) on the Atlassian website.

- **Optional:** Configured OAuth 2.0 credentials containing a Confluence app key, Confluence app secret, Confluence access token, and Confluence refresh token to allow Amazon Q to connect to your Confluence instance. If your access token expires, you can either use the refresh token to regenerate your access token and refresh token pair. Or, you can repeat the authorization process.
- **Optional:** Configured a Personal Access Token (PAT) containing a Confluence token to allow Amazon Q to connect to your Confluence Server/Data Center instance. For information about how to create a PAT token, see [Using Personal Access Tokens](#) on the Atlassian website.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your Confluence (Server/Data Center) authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Using the console

On the **Confluence** page, enter the following information:

1. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

2. In **Source**, enter the following information:

- a. In **Source**, for **Hosting Method** – Choose **Confluence Server/Data Center**.

- b. **Confluence URL** – Enter the Confluence host URLs. The format for the host URL that you enter is *https://example.confluence.com*.

 **Important**

If you change or update your Confluence (Server/Data Center) data source URL, you also need to update your Secrets Manager secret to ensure a secure connection.

- c. **SSL certificate location** – Enter the file path to an SSL certificate stored in an Amazon S3 bucket.
3. **Web proxy** – *optional*, enter the following information:
 - a. **Host name** – Host name for your Confluence account.
 - b. **Port number** – Port used by the host URL transport protocol.
 4. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

 **Note**

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

5. For **Authentication** – Choose between **Basic authentication**, **Oauth 2.0 authentication**, and **Personal Access Token authentication** based on your use case.
6. **AWS Secrets Manager secret** – Choose an existing secret or create a Secrets Manager secret to store your Confluence authentication credentials. If you choose to create a secret, an AWS Secrets Manager secret window opens. Enter the following information in the window:
 - a. **Secret name** – A name for your secret.
 - b. If using **Basic Authentication** – Enter the **Secret name Username**, and **Password** (Confluence Server/Data Center password) that you generated and downloaded from your Confluence account.

If using **OAuth2.0 Authentication** – Enter the **Secret name**, **App key**, **App secret**, **Access token**, and **Refresh token** you created in your Confluence account.

If using **Personal Access Token authentication** – Enter the **Secret name** and the **Confluence Server PAT token** that you created in your Confluence Server account.

- c. Choose **Save and add secret**.
7. **Configure VPC and security group – optional** – Choose whether you want to use a VPC. If you do, enter the following information:
 - a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
 - b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

8. **Identity crawler** – Choose to activate Amazon Q identity crawler to sync identity information. For more information, see [Identity crawler](#).
9. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

10. In **Sync scope**, choose from the following options:
 - a. In **Sync scope**, for **sync contents**, choose to sync from the following entity types: **Pages**, **Page comments**, **Page attachments**, **Blogs**, **Blog comments**, **Blog attachments**, **Personal spaces**, and **Archived spaces**.

 **Note**

Page comments and **Page attachments** can only be selected if you choose to sync **Pages**. **Blog comments** and **Blog attachments** can only be selected if you choose to sync **Blogs**.

⚠ Important

You can crawl **Pages** and **Blogs** from one of more specific **Spaces**. If you don't specify a **Space key** regex pattern in **Additional configuration**, all **Pages** and **Blogs** will be crawled by default. If no **Space** is specified in the filter, all spaces will be crawled.

b. In **Additional configuration – optional**, for **Space and regex patterns**, specify whether to include or exclude specific spaces in your index using:

- **Space key** – For example, *my-space-123*.

ℹ Note

If you don't specify a **Space key** regex pattern in **Additional configuration**, all **Pages** and **Blogs** will be crawled by default. If no **Space** is specified in the filter, all spaces will be crawled.

- **URL** – For example, *.*MySite/MyDocuments/*.
- **File type** – For example, *.*\pdf*, *.*\txt*.
- For **Maximum file size** – Specify the file size limit in MBs that Amazon Q will crawl. Amazon Q will crawl only the files within the size limit you define. The default file size is 50MB. The maximum file size should be greater than 0MB and less than or equal to 50MB.
- For **Entity title regex patterns** – Specify regular expression patterns to include or exclude certain **Blogs**, **Pages**, **Comments**, and **Attachments** by titles.

ℹ Note

If you want to crawl a specific page or subpage, you can use page title regex patterns to either include or exclude this page.

11. For **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.
 - **Full sync** – Sync all content regardless of the previous sync status.
 - **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.
12. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
13. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
14. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:
 - a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
 - b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

 **Note**

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

15. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

16. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

Note

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to Confluence (Server/Data Center) using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the `configuration` parameter to provide a JSON schema with all other configuration information specific to your data source connector.

For an example of the API request, see [CreateDataSource](#) in the Amazon Q API Reference.

Confluence JSON schema

The following is the Confluence JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
      "type": "object",
      "properties": {
        "repositoryEndpointMetadata": {
          "type": "object",
          "properties": {
            "hostUrl": {
              "type": "string",
              "pattern": "https:.*"
            },
            "type": {
              "type": "string",
              "enum": [
                "SAAS",
                "ON_PREM"
              ]
            }
          }
        },
      }
    },
  },
}
```

```
    "authType": {
      "type": "string",
      "enum": [
        "Basic",
        "OAuth2",
        "Personal-token"
      ]
    },
    "required": [
      "hostUrl",
      "type",
      "authType"
    ]
  },
  "required": [
    "repositoryEndpointMetadata"
  ],
  "repositoryConfigurations": {
    "type": "object",
    "properties": {
      "space": {
        "type": "object",
        "properties": {
          "fieldMappings": {
            "type": "array",
            "items": [
              {
                "type": "object",
                "properties": {
                  "indexFieldName": {
                    "type": "string"
                  },
                  "indexFieldType": {
                    "type": "string",
                    "enum": [
                      "STRING",
                      "STRING_LIST",
                      "DATE"
                    ]
                  }
                }
              }
            ]
          },
          "dataSourceFieldName": {
```

```

        "type": "string"
      },
      "dateFieldFormat": {
        "type": "string",
        "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
      }
    },
    "required": [
      "indexFieldName",
      "indexFieldType",
      "dataSourceFieldName"
    ]
  }
]
}
},
"required": [
  "fieldMappings"
]
},
"page": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": [
        {
          "type": "object",
          "properties": {
            "indexFieldName": {
              "type": "string"
            },
            "indexFieldType": {
              "type": "string",
              "enum": [
                "STRING",
                "STRING_LIST",
                "DATE",
                "LONG"
              ]
            },
            "dataSourceFieldName": {
              "type": "string"
            }
          }
        }
      ]
    }
  }
}

```

```

        "dateFieldFormat": {
            "type": "string",
            "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
        }
    },
    "required": [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
    ]
}
]
}
},
"required": [
    "fieldMappings"
]
},
"blog": {
    "type": "object",
    "properties": {
        "fieldMappings": {
            "type": "array",
            "items": [
                {
                    "type": "object",
                    "properties": {
                        "indexFieldName": {
                            "type": "string"
                        },
                        "indexFieldType": {
                            "type": "string",
                            "enum": [
                                "STRING",
                                "STRING_LIST",
                                "DATE",
                                "LONG"
                            ]
                        },
                        "dataSourceFieldName": {
                            "type": "string"
                        },
                        "dateFieldFormat": {
                            "type": "string",

```



```

        "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
    }
},
"required": [
    "indexFieldName",
    "indexFieldType",
    "dataSourceFieldName"
]
}
]
}
},
"required": [
    "fieldMappings"
]
},
"comment": {
    "type": "object",
    "properties": {
        "fieldMappings": {
            "type": "array",
            "items": [
                {
                    "type": "object",
                    "properties": {
                        "indexFieldName": {
                            "type": "string"
                        },
                        "indexFieldType": {
                            "type": "string",
                            "enum": [
                                "STRING",
                                "STRING_LIST",
                                "DATE",
                                "LONG"
                            ]
                        },
                        "dataSourceFieldName": {
                            "type": "string"
                        },
                        "dateFieldFormat": {
                            "type": "string",
                            "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
                        }
                    }
                }
            ]
        }
    }
}

```

```

        },
        "required": [
            "indexFieldName",
            "indexFieldType",
            "dataSourceFieldName"
        ]
    }
]
}
},
"required": [
    "fieldMappings"
]
},
"attachment": {
    "type": "object",
    "properties": {
        "fieldMappings": {
            "type": "array",
            "items": [
                {
                    "type": "object",
                    "properties": {
                        "indexFieldName": {
                            "type": "string"
                        },
                        "indexFieldType": {
                            "type": "string",
                            "enum": [
                                "STRING",
                                "STRING_LIST",
                                "DATE",
                                "LONG"
                            ]
                        },
                        "dataSourceFieldName": {
                            "type": "string"
                        },
                        "dateFieldFormat": {
                            "type": "string",
                            "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
                        }
                    }
                }
            ]
        },
        "required": [

```

```

        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
    ]
    }
    ]
}
},
"required": [
    "fieldMappings"
]
}
}
},
"additionalProperties": {
    "type": "object",
    "properties": {
        "isCrawlAcl": {
            "type": "boolean"
        },
        "fieldForUserId": {
            "type": "string"
        },
        "inclusionSpaceKeyFilter": {
            "type": "array",
            "items": {
                "type": "string"
            }
        },
        "exclusionSpaceKeyFilter": {
            "type": "array",
            "items": {
                "type": "string"
            }
        },
        "pageTitleRegEX": {
            "type": "array",
            "items": {
                "type": "string"
            }
        },
        "blogTitleRegEX": {
            "type": "array",
            "items": {

```

```
    "type": "string"
  }
},
"commentTitleRegEX": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"attachmentTitleRegEX": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"isCrawlPersonalSpace": {
  "type": "boolean"
},
"isCrawlArchivedSpace": {
  "type": "boolean"
},
"isCrawlArchivedPage": {
  "type": "boolean"
},
"isCrawlPage": {
  "type": "boolean"
},
"isCrawlBlog": {
  "type": "boolean"
},
"isCrawlPageComment": {
  "type": "boolean"
},
"isCrawlPageAttachment": {
  "type": "boolean"
},
"isCrawlBlogComment": {
  "type": "boolean"
},
"isCrawlBlogAttachment": {
  "type": "boolean"
},
"maxFileSizeInMegaBytes": {
  "type": "string"
}
```

```
    },
    "inclusionFileTypePatterns": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    "exclusionFileTypePatterns": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    "inclusionUrlPatterns": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    "exclusionUrlPatterns": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    "proxyHost": {
      "type": "string"
    },
    "proxyPort": {
      "type": "string"
    }
  },
  "required": [],
  "type": {
    "type": "string",
    "enum": [
      "CONFLUENCEV2",
      "CONFLUENCE"
    ]
  },
  "enableIdentityCrawler": {
    "type": "boolean"
  },
}
```

```

"syncMode": {
  "type": "string",
  "enum": [
    "FULL_CRAWL",
    "FORCED_FULL_CRAWL"
  ]
},
"secretArn": {
  "type": "string",
  "minLength": 20,
  "maxLength": 2048
}
},
"version": {
  "type": "string",
  "anyOf": [
    {
      "pattern": "1.0.0"
    }
  ]
},
"required": [
  "connectionConfiguration",
  "repositoryConfigurations",
  "syncMode",
  "additionalProperties",
  "secretArn",
  "type"
]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.
repositoryEndpointMetadata	The endpoint information for the data source.

Configuration	Description
hostUrl	<p>The URL for your Confluence instance. For example, <i>https://example.confluence.com</i> .</p> <div data-bbox="829 401 1507 762" style="border: 1px solid #f08080; border-radius: 10px; padding: 10px; background-color: #fff9e6;"> <p>⚠ Important</p> <p>If you change or update your Confluence (Server/Data Center) data source URL, you also need to update your Secrets Manager secret to ensure a secure connection.</p> </div>
type	The hosting method for your Confluence instance, whether SAAS or ON_PREM.
authType	The authentication method for your Confluence instance, whether Basic, OAuth2, or Personal-token .
repositoryConfigurations <ul style="list-style-type: none"> • space • page • blog • comment • attachment 	<p>Configuration information for the content of the data source. For example, configuring specific types of content and field mappings.</p> <p>A list of objects that map the attributes or field names of your Confluence spaces, pages, blogs, comments, and attachments to Amazon Q index field names.</p>
additionalProperties	Additional configuration options for your content in your data source.
isCrawlAcl	Specify true to crawl access control information from documents.

Configuration	Description
fieldForUserId	Specify field to use for UserId for ACL crawling.
proxyHost	The host where the web proxy is required. The host name should be without protocol (http:// or https://).
proxyPort	Port used by the host URL transport protocol. The port number should be a numeric value between 0 and 65535.
maxFileSizeInMegabytes	Specify the file size limit in MBs that Amazon Q will crawl. Amazon Q will crawl only the files within the size limit you define. The default file size is 50MB. The maximum file size should be greater than 0MB and less than or equal to 50MB.
<ul style="list-style-type: none"> • inclusionSpaceKeyFilter • exclusionSpaceKeyFilter • pageTitleRegEX • blogTitleRegEX • commentTitleRegEX • attachmentTitleRegEX • inclusionFileTypePatterns • exclusionFileTypePatterns • inclusionUrlPatterns • exclusionUrlPatterns 	A list of regular expression patterns to include and/or exclude certain files in your Confluence data source. Files that match the patterns are included in the index. Files that don't match the patterns are excluded from the index. If a file matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence and the file isn't included in the index.

Configuration	Description
<ul style="list-style-type: none"> • <code>isCrawlPersonalSpace</code> • <code>isCrawlArchivedSpace</code> • <code>isCrawlArchivedPage</code> • <code>isCrawlPage</code> • <code>isCrawlBlog</code> • <code>isCrawlPageComment</code> • <code>isCrawlPageAttachment</code> • <code>isCrawlBlogComment</code> • <code>isCrawlBlogAttachment</code> 	<p><code>true</code> to index files in your Confluence personal spaces, pages, blogs, page comments, page attachments, blog comments, and blog attachments.</p>
<p><code>type</code></p>	<p>The type of data source. Specify <code>CONFLUENCE_V2</code> as your data source type.</p>
<p><code>enableIdentityCrawler</code></p>	<p><code>true</code> to activate identity crawler. Identity crawler is activated by default. See Identity crawler for more information.</p>
<p><code>syncMode</code></p>	<p>Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose between the following options:</p> <ul style="list-style-type: none"> • Use <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index • Use <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index

Configuration	Description
secretARN	<p>The Amazon Resource Name (ARN) of a Secrets Manager secret that contains the key-value pairs required to connect to your Confluence instance.</p> <p>If you use OAuth 2.0 authentication, the secret must contain a JSON structure with the following keys:</p> <pre data-bbox="829 615 1507 1171">{ "hostUrl": " <i>Confluence Server host URL</i> ", "confluenceAppKey": " <i>client ID for your Confluence account</i> ", "confluenceAppSecret": " <i>client secret from your Confluence token</i> ", "confluenceAccessToken": " <i>access token created in Confluence</i> ", "confluenceRefreshToken": " <i>refresh token created in Confluence</i> " }</pre> <p>(For Confluence Server/Data Center only) If you use basic authentication, the secret is stored in a JSON structure with the following keys:</p> <pre data-bbox="829 1419 1507 1772">{ "hostUrl": " <i>Confluence Server/Data Center host URL</i> ", "username": " <i>Confluence Server/Data Center user name</i> ", "password": " <i>Confluence Server/Data Center password</i> " }</pre>

Configuration	Description
	<p>(For Confluence Server/Data Center only) If you use Personal Access Token authentication, the secret is stored in a JSON structure with the following keys:</p> <pre data-bbox="829 426 1507 663"> { "hostUrl": " <i>Confluence Server/</i> <i>Data Center host URL</i> ", "patToken": " <i>Confluence token</i> " } </pre>
version	The version of this template that's currently supported.

How Amazon Q Business connector crawls Confluence (Server/Data Center) ACLs

When you connect an Confluence (Server/Data Center) data source to Amazon Q Business, Amazon Q crawls ACL information attached to a document (user and group information) from your Confluence (Server/Data Center) instance. If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

You configure user and group access to spaces using the space permissions page. For pages and blogs, you use the restrictions page. For more information about space permissions, see [Space Permissions Overview](#) on the Confluence Support website. For more information about page and blog restrictions, see [Page Restrictions](#) on the Confluence Support website.

The group and user IDs are mapped as follows:

- `_group_ids` – Group names are present on spaces, pages, and blogs where there are restrictions. They're mapped from the name of the group in Confluence . Group names are always lower case.
- `_user_id` – User names are present on the space, page, or blog where there are restrictions. They're mapped depending on the type of Confluence instance that you are using.
- For Confluence Cloud – The `_user_id` is the account ID of the user.

⚠ Important

For user context filtering to work correctly for your Confluence connector, you need to make sure that the visibility of a user granted access to a Confluence page is set to **Anyone**. For more information, see [Set your email visibility](#) in Atlassian Developer Documentation.

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business Confluence (Server/Data Center) data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q Business enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have an attribute mapping already available, or if you want to map additional document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

⚠ Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q Confluence connector supports the following entities and the associated reserved and custom attributes.

Supported entities and field mappings

- [Space](#)
- [Page](#)
- [Blog](#)
- [Comment](#)
- [Attachment](#)

Space

Confluence field name	Index field name	Description	Data type
spaceName	cf_sp_document_title	Custom	String
itemType	_category	Default	String
url	_source_uri	Default	String
spaceKey	cf_space_key	Custom	String
description	cf_description	Custom	String
spaceType	cf_type	Custom	String

Page

Confluence field name	Index field name	Description	Data type
title	_cf_page_document_title	Custom	String
authors	_authors	Default	String list
createdDate	_created_at	Default	Date
modifiedDate	_last_updated_at	Default	Date
labels	cf_labels	Custom	String list
version	cf_version	Custom	Long (numeric)
itemType	_category	Default	String
spaceKey	cf_space_key	Custom	String
spaceName	cf_space_name	Custom	String
url	_source_uri	Default	String
status	cf_status	Custom	String
parentId	cf_parent_id	Custom	String

Blog

Confluence field name	Index field name	Description	Data type
title	cf_bg_document_title	Custom	String
author	_authors	Default	String list
publishedDate	_created_at	Default	Date

Confluence field name	Index field name	Description	Data type
labels	_source_uri	Default	String
version	cf_version	Custom	Long (numeric)
itemType	_category	Custom	String
spaceKey	cf_space_key	Custom	String
modifiedDate	_last_updated_at	Default	Date
spaceName	cf_space_name	Custom	String
status	cf_status	Custom	String
url	_source_uri	Default	String
parentId	cf_parent_id	Custom	String

Comment

Confluence field name	Index field name	Description	Data type
title	cf_cmt_document_title	Custom	String
author	_authors	Default	String list
createdDate	_created_at	Default	Date
version	cf_version	Custom	Long (numeric)
itemType	_category	Default	String
spaceKey	cf_space_key	Custom	String
spaceName	cf_space_name	Custom	String

Confluence field name	Index field name	Description	Data type
contentType	cf_content_type	Custom	String
url	_source_uri	Default	String
parentId	cf_parent_id	Custom	String
status	cf_status	Custom	String

Attachment

Confluence field name	Index field name	Description	Data type
fileName	cf_attachment_document_title	Custom	String
author	_authors	Default	String list
createdDate	_created_at	Default	Date
labels	cf_labels	Custom	String list
version	cf_version	Custom	Long (numeric)
itemType	_category	Default	String
spaceKey	cf_space_key	Custom	String
contentType	cf_content_type	Custom	String
modifiedDate	_last_updated_at	Default	Date
fileSize	cf_file_size	Custom	Long (numeric)
fileType	cf_attachment_file_type	Custom	String

Confluence field name	Index field name	Description	Data type
spaceName	cf_space_name	Custom	String
documentId	_document_id	Default	String list
url	_source_uri	Default	String
parentId	cf_parent_id	Custom	String
attachmentComment	cf_attachment_comment	Custom	String
status	cf_status	Custom	String

IAM role for Amazon Q Confluence (Server/Data Center) connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the `BatchPutDocument` and `BatchDeleteDocument` operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- Permission to access the SSL certificate stored in your Amazon S3 bucket.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```

{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AllowsAmazonQToGetS3Objects",
    "Action": [
      "s3:GetObject"
    ],
    "Resource": [
      "arn:aws:s3:::{{input_bucket_name}}/*"
    ],
    "Effect": "Allow",
    "Condition": {
      "StringEquals": {
        "aws:ResourceAccount": "{{account_id}}"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToGetSecret",
    "Effect": "Allow",
    "Action": [
      "secretsmanager:GetSecretValue"
    ],
    "Resource": [
      "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToDecryptSecret",
    "Effect": "Allow",
    "Action": [
      "kms:Decrypt"
    ],
    "Resource": [
      "arn:aws:kms:{{region}}:{{account_id}}:key/[key_id]"
    ],
    "Condition": {
      "StringLike": {
        "kms:ViaService": [
          "secretsmanager.*.amazonaws.com"
        ]
      }
    }
  }
}

```

```

    },
    {
      "Sid": "AllowsAmazonQToIngestDocuments",
      "Effect": "Allow",
      "Action": [
        "qbusiness:BatchPutDocument",
        "qbusiness:BatchDeleteDocument"
      ],
      "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
    },
    {
      "Sid": "AllowsAmazonQToIngestPrincipalMapping",
      "Effect": "Allow",
      "Action": [
        "qbusiness:PutGroup",
        "qbusiness:CreateUser",
        "qbusiness>DeleteGroup",
        "qbusiness:UpdateUser",
        "qbusiness:ListGroups"
      ],
      "Resource": [
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}",
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}",
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}/data-source/*"
      ]
    },
    {
      "Sid": "AllowsAmazonQToCreateAndDeleteNI",
      "Effect": "Allow",
      "Action": [
        "ec2:CreateNetworkInterface",
        "ec2>DeleteNetworkInterface"
      ],
      "Resource": [
        "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
        "arn:aws:ec2:{{region}}:{{account_id}}:security-group/
[{{security_group}}]"
      ]
    },
    {

```

```

    "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
    "Effect": "Allow",
    "Action": [
        "ec2:CreateNetworkInterface",
        "ec2:DeleteNetworkInterface"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
        "StringLike": {
            "aws:RequestTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
        },
        "ForAllValues:StringEquals": {
            "aws:TagKeys": [
                "AMAZON_Q"
            ]
        }
    }
},
{
    "Sid": "AllowsAmazonQToCreateTags",
    "Effect": "Allow",
    "Action": [
        "ec2:CreateTags"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
        "StringEquals": {
            "ec2:CreateAction": "CreateNetworkInterface"
        }
    }
},
{
    "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
    "Effect": "Allow",
    "Action": [
        "ec2:CreateNetworkInterfacePermission"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
        "StringLike": {
            "aws:ResourceTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
        }
    }
}

```

```

    }
  },
  {
    "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
    "Effect": "Allow",
    "Action": [
      "ec2:DescribeNetworkInterfaces",
      "ec2:DescribeAvailabilityZones",
      "ec2:DescribeNetworkInterfaceAttribute",
      "ec2:DescribeVpcs",
      "ec2:DescribeRegions",
      "ec2:DescribeNetworkInterfacePermissions",
      "ec2:DescribeSubnets"
    ],
    "Resource": "*"
  }
]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToAssumeRoleForServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "{{source_account}}"
        },
        "ArnLike": {
          "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
        }
      }
    }
  ]
}

```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Troubleshooting your Amazon Q Business Confluence (Server/Data Center) connector

The following table provides information about error codes you may see for the Confluence (Server/Data Center) connector and suggested troubleshooting actions.

Error code	Error message	Suggested resolution
CNF-5500	Null/empty username.	Provide username.
CNF-5501	Error validating credentials due to Invalid username or password.	Provide valid username/password
CNF-5502	Null/empty confluence AppKey.	Provide confluence AppKey.
CNF-5503	Null/empty confluence Secret.	Provide confluence Secret.
CNF-5504	Null/empty Client Access Token.	Provide Client Access Token.
CNF-5505	Null/empty Client Refresh Token	Provide Client Refresh Token
CNF-5506	Incorrect auth type.	Auth type should be Basic or OAuth2 or Personal-token.
CNF-5507	Null/empty auth type.	Auth Type should not be null or empty value.
CNF-5508	Empty/null host URL.	Host Url should not be null or empty value.
CNF-5509	Null/empty crawl type.	Crawl Type should not be null or empty value.

Error code	Error message	Suggested resolution
CNF-5510	Null/empty Repository Configurations.	Repository Configurations should not be null or empty value.
CNF-5511	Incorrect type.	type should be SAAS or ON_PREM.
CNF-5512	Invalid inclusion file type patterns.	Provide the correct inclusion patterns.
CNF-5513	Invalid exclusion file type patterns.	Provide the correct exclusion patterns.
CNF-5514	Invalid regex patterns.	Provide the correct regex patterns.
CNF-5515	Error validating credentials due to invalid username or password.	Provide valid user name and password.
CNF-5516	Error validating credentials due to invalid client id or client secret.	Provide valid client id and client secret.
CNF-5517	Error validating crawl type.	Provide valid crawl type.
CNF-5518	Invalid URI.	Provide valid URI.
CNF-5519	Null/empty DataSourceFieldName in Space Entity.	Provide value for DataSourceFieldName in Space Entity.
CNF-5520	Null/empty IndexFieldName in Blog Entity.	Provide value for IndexFieldName in Blog Entity.
CNF-5521	Null/empty IndexFieldType in Space Entity.	Provide value for IndexFieldType in Space Entity.
CNF-5522	Null/empty password.	Provide password.

Error code	Error message	Suggested resolution
CNF-5523	Incorrect auth type.	Auth type should be Basic or OAuth2.
CNF-5524	Null/empty DataSourceFieldName in Page Entity.	Provide value for DataSourceFieldName in Page Entity.
CNF-5525	Null/empty DataSourceFieldName in Blog Entity	Please provide value for DataSourceFieldName in Blog Entity
CNF-5526	Null/empty DataSourceFieldName in Comment Entity.	Provide value for DataSourceFieldName in Comment Entity.
CNF-5527	Null/empty DataSourceFieldName in Attachment Entity.	Provide value for DataSourceFieldName in Attachment Entity.
CNF-5528	Null/empty IndexFieldName.	IndexFieldName field can't be null or empty value.
CNF-5529	Null/empty IndexFieldName in Space Entity.	Provide value for IndexFieldName in Space Entity.
CNF-5530	Null/empty IndexFieldName in Page Entity	Please provide value for IndexFieldName in Page Entity
CNF-5531	Invalid isCrawlPersonalSpace value.	isCrawlPersonalSpace should be a boolean value true or false.
CNF-5532	Invalid isCrawlArchivedSpace value.	isCrawlArchivedSpace should be a boolean value true or false.
CNF-5533	Invalid isCrawlArchivedPage value.	isCrawlArchivedPage should be a boolean value true or false.

Error code	Error message	Suggested resolution
CNF-5534	Invalid isCrawlPage value.	isCrawlPage should be a boolean value true or false.
CNF-5535	Invalid isCrawlBlogComment value.	isCrawlBlogComment should be a boolean value true or false.
CNF-5536	Invalid isCrawlBlogComment value.	isCrawlBlogComment should be a boolean value true or false.
CNF-5537	Invalid isCrawlBlogAttachment value.	isCrawlBlogAttachment should be a boolean value true or false.
CNF-5538	Error validating on protocol.	Provide valid protocol.
CNF-5539	Null/empty IndexFieldName in Comment Entity.	Provide value for IndexFieldName in Comment Entity.
CNF-5540	Null/empty Personal Access Token.	Provide Personal Access Token.
CNF-5541	Invalid OAuth value.	Give a valid OAuth URL.
CNF-5542	Invalid Space value.	Give a valid Space URL.
CNF-5543	Archived Space Exception .	Check Archived Space.
CNF-5544	JSON Exception for Space.	Check Space.
CNF-5545	JSON Exception for Comment.	Check Comment.
CNF-5546	JSON Exception for Comment.	Check Comment.

Error code	Error message	Suggested resolution
CNF-5547	JSON Exception for Comment.	Check Comment.
CNF-5548	JSON Exception for Attachment.	Check Attachment.
CNF-5549	JSON Exception for Blog.	Check Blog.
CNF-5550	JSON Exception for Page.	Check Page.
CNF-5551	JSON Exception for Label.	Check Label.
CNF-5552	JSON Exception for ACL.	Check ACL.
CNF-5553	JSON Exception for Groups.	Check Groups.
CNF-5554	JSON Exception for Group Members.	Check Group Members.
CNF-5555	JSON Exception for Space Group.	Check Space Group.
CNF-5556	Exception in CommentItem.	Check the CommentItem class.
CNF-5557	Invalid isCrawlPageComment value.	isCrawlPageComment should be a boolean value true or false.
CNF-5558	Invalid isCrawlPageAttachment value.	isCrawlPageAttachment should be a boolean value true or false.
CNF-5559	Null/empty Repository Configurations.	Repository Configurations should not be null or empty value.
CNF-5560	Null/empty IndexFieldName in Attachment.	Please provide value for IndexFieldName in Attachment Entity.

Error code	Error message	Suggested resolution
CNF-5561	Invalid proxy url.	Proxy url should not contain http: or https.
CNF-5562	Null/Empty proxy port.	Provide a valid proxy port.
CNF-5563	Invalid Host URL.	Provide valid Host URL.
CNF-5564	Invalid proxy port value.	Provide a valid proxy port.
CNF-5565	Confluence server not reachable.	Provide a valid proxy and server details.
CNF-5566	Null/empty IndexField dType in Page Entity.	Provide value for IndexFieldType in Page Entity.
CNF-5567	Null/empty IndexField dType in Blog Entity.	Provide value for IndexFieldType in Blog Entity.
CNF-5568	Null/empty IndexField dType in Comment Entity.	Provide value for IndexFieldType in Comment Entity.
CNF-5569	Null/empty IndexField dType in Attachment. Entity	Provide value for IndexFieldType in Attachment. Entity
CNF-5570	JSON Exception for Content Ancestors.	Check your Ancestors.
CNF-5571	Invalid Host URL Pattern.	Provide valid Host URL Pattern.
CNF-5572	Error validating credentials due to Invalid access or refresh token.	Invalid AccessToken/RefreshToken.

Connecting Dropbox to Amazon Q Business

Dropbox is a file hosting service that offers cloud storage, document organization, and document templating services. You can connect Dropbox instance to Amazon Q Business—using either the AWS Management Console or the [CreateDataSource](#) API—and create an Amazon Q web experience.

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).
- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [Dropbox connector overview](#)
- [Prerequisites for connecting Amazon Q Business to Dropbox](#)
- [Connecting Amazon Q Business to Dropbox using the console](#)
- [Connecting Amazon Q Business to Dropbox using APIs](#)
- [How Amazon Q Business connector crawls Dropbox ACLs](#)
- [Amazon Q Business Dropbox data source connector field mappings](#)
- [IAM role for Amazon Q Business Dropbox connector](#)

Dropbox connector overview

The following table gives an overview of the Amazon Q Business Dropbox connector and its supported features.

Category	Feature	Support
Security	Authentication type	Permanent token (recommended), Access token (temporary use)
	Authentication credentials	Permanent token <ul style="list-style-type: none"> • App key

Category	Feature	Support
		<ul style="list-style-type: none"> App secret Permanent token <p>Access token</p> <ul style="list-style-type: none"> App key App secret Access token
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Identity crawling	Yes
	VPC	Yes
Crawl features	Custom metadata	Yes
	Entities	<p>Yes. The following entities are supported:</p> <ul style="list-style-type: none"> Files Dropbox Paper Dropbox Paper Templates Shortcuts
	Field mappings	Yes. Supports default and custom field mappings. For more information, see Field mappings .
	Filters	<p>Yes. The following filters are supported:</p> <ul style="list-style-type: none"> Include/ exclude Files Dropbox Paper, Dropbox Paper templates, and Shortcuts. Include/exclude content by file name, file type, and file path.

Category	Feature	Support
	Sync mode	Supports full and incremental sync.
	File types	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to Dropbox

Before you begin, make sure that you have completed the following prerequisites.

In Dropbox, make sure you have:

- Created a Dropbox Advanced account and set up an admin user.
- Created a Dropbox app with a unique **App name**, activated **Scoped Access**. For more information, see [Dropbox documentation on creating an app](#) on the Dropbox website.
- Activated **Full Dropbox** permissions on the Dropbox console and added the following permissions:
 - files.content.read
 - files.metadata.read
 - sharing.read
 - file_requests.read
 - groups.read
 - team_info.read
 - team_data.content.read
- Noted your Dropbox app key, Dropbox app secret, and Dropbox access token for basic authentication credentials.
- Generated and copied a temporary OAuth 2.0 access token for your Dropbox app. This token is temporary and expires after 4 hours. For more information, see [Dropbox documentation on OAuth authentication](#) on the Dropbox website.

Recommended: Configured a Dropbox permanent refresh token that never expires to allow Amazon Q to continue to sync your data source without any disruptions. For more information, see [Dropbox documentation on refresh tokens](#) on the Dropbox website.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your Dropbox authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Connecting Amazon Q Business to Dropbox using the console

The following procedure outlines how to connect Amazon Q Business to Dropbox using the AWS Management Console.

Connecting Amazon Q to Dropbox

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **Dropbox** page, enter the following information:
6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

Note

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

8. In **Authentication** – Choose between **Permanent Token (recommended)** and **Access Token (temporary use)** based on your use case.
9. In **Authentication credentials**, for **AWS Secrets Manager secret** – Choose an existing secret or create a Secrets Manager secret to store your Dropbox authentication credentials. If you choose to create a secret, an AWS Secrets Manager secret window opens.
 - Enter following information in the **Create an AWS Secrets Manager secret window**:
 - i. **Secret name** – A name for your secret.
 - ii. For **App key**, **App secret**, and token information (permanent or temporary) – Enter the authentication credential values that you generated from your Dropbox account.
 - iii. Choose **Save**.
10. **Configure VPC and security group** – *optional* – Choose whether you want to use a VPC. If you do, enter the following information:
 - a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
 - b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

11. **Identity crawler** – Choose to activate Amazon Q identity crawler to sync identity information. For more information, see [Identity crawler](#).
12. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

13. In **Sync scope**, enter the following information.

- a. For **Select entities or content types** – Choose entities or content types you want to crawl.
 - b. **Change log mode** – Choose to update your index instead of syncing all files.
 - c. In **Additional configuration – optional**, for **Regex patterns** – Add regular expression patterns to include or exclude certain files.
14. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
15. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
16. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:
- a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
 - b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

 **Note**

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

17. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

18. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

 **Note**

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to

view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to Dropbox using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the configuration parameter to provide a JSON schema with all other configuration information specific to your data source connector.

For an example of the API request, see [CreateDataSource](#) in the Amazon Q API Reference.

Dropbox JSON schema

The following is the Dropbox JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
      "type": "object",
      "properties": {
        "repositoryEndpointMetadata": {
          "type": "object",
          "properties": {
          }
        }
      }
    },
    "required": [
      "repositoryEndpointMetadata"
    ]
  },
  "repositoryConfigurations": {
    "type": "object",
    "properties": {
      "file": {
        "type": "object",
        "properties": {
          "fieldMappings": {
            "type": "array",
            "items": {

```

```

    "anyOf": [
      {
        "type": "object",
        "properties": {
          "indexFieldName": {
            "type": "string"
          },
          "indexFieldType": {
            "type": "string",
            "enum": [
              "STRING",
              "STRING_LIST",
              "LONG",
              "DATE"
            ]
          },
          "dataSourceFieldName": {
            "type": "string"
          },
          "dateFieldFormat": {
            "type": "string",
            "pattern": "dd-MM-yyyy HH:mm:ss"
          }
        },
        "required": [
          "indexFieldName",
          "indexFieldType",
          "dataSourceFieldName"
        ]
      }
    ],
    "required": [
      "fieldMappings"
    ],
    "paper": {
      "type": "object",
      "properties": {
        "fieldMappings": {
          "type": "array",
          "items": {

```

```

    "anyOf": [
      {
        "type": "object",
        "properties": {
          "indexFieldName": {
            "type": "string"
          },
          "indexFieldType": {
            "type": "string",
            "enum": [
              "STRING",
              "STRING_LIST",
              "LONG",
              "DATE"
            ]
          },
          "dataSourceFieldName": {
            "type": "string"
          },
          "dateFieldFormat": {
            "type": "string",
            "pattern": "dd-MM-yyyy HH:mm:ss"
          }
        },
        "required": [
          "indexFieldName",
          "indexFieldType",
          "dataSourceFieldName"
        ]
      }
    ],
    "required": [
      "fieldMappings"
    ],
    "papert": {
      "type": "object",
      "properties": {
        "fieldMappings": {
          "type": "array",
          "items": {

```

```

    "anyOf": [
      {
        "type": "object",
        "properties": {
          "indexFieldName": {
            "type": "string"
          },
          "indexFieldType": {
            "type": "string",
            "enum": [
              "STRING",
              "STRING_LIST",
              "LONG",
              "DATE"
            ]
          },
          "dataSourceFieldName": {
            "type": "string"
          },
          "dateFieldFormat": {
            "type": "string",
            "pattern": "dd-MM-yyyy HH:mm:ss"
          }
        },
        "required": [
          "indexFieldName",
          "indexFieldType",
          "dataSourceFieldName"
        ]
      }
    ],
    "required": [
      "fieldMappings"
    ],
    "shortcut": {
      "type": "object",
      "properties": {
        "fieldMappings": {
          "type": "array",
          "items": {

```

```
    "anyOf": [
      {
        "type": "object",
        "properties": {
          "indexFieldName": {
            "type": "string"
          },
          "indexFieldType": {
            "type": "string",
            "enum": [
              "STRING",
              "STRING_LIST",
              "LONG",
              "DATE"
            ]
          },
          "dataSourceFieldName": {
            "type": "string"
          },
          "dateFieldFormat": {
            "type": "string",
            "pattern": "dd-MM-yyyy HH:mm:ss"
          }
        },
        "required": [
          "indexFieldName",
          "indexFieldType",
          "dataSourceFieldName"
        ]
      }
    ],
    "required": [
      "fieldMappings"
    ]
  },
  "secretArn": {
    "type": "string"
  },
  "enableIdentityCrawler": {
```

```
"type": "boolean"
},
"additionalProperties": {
  "type": "object",
  "properties": {
    "isCrawlAcl": {
      "type": "boolean"
    },
    "fieldForUserId": {
      "type": "string"
    },
    "inclusionPatterns": {
      "type": "array"
    },
    "exclusionPatterns": {
      "type": "array"
    },
    "crawlFile": {
      "type": "boolean"
    },
    "crawlPaper": {
      "type": "boolean"
    },
    "crawlPapert": {
      "type": "boolean"
    },
    "crawlShortcut": {
      "type": "boolean"
    }
  }
},
"type": {
  "type": "string",
  "pattern": "DROPBOX"
},
"useChangeLog": {
  "type": "string",
  "enum": [
    "true",
    "false"
  ]
},
"tokenType": {
  "type": "string",
```

```

    "enum": [
      "PERMANENT",
      "TEMPORARY"
    ]
  },
  "version": {
    "type": "string",
    "anyOf": [
      {
        "pattern": "1.0.0"
      }
    ]
  }
},
"additionalProperties": false,
"required": [
  "connectionConfiguration",
  "repositoryConfigurations",
  "additionalProperties",
  "useChangeLog",
  "secretArn",
  "type",
  "tokenType"
]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.
repositoryEndpointMetadata	The endpoint information for the data source. This data source doesn't specify an endpoint in <code>repositoryEndpointMetadata</code> . Rather, the connection information is included in an AWS Secrets Manager secret that you provide the <code>secretArn</code> .

Configuration	Description
repositoryConfigurations	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings.
<ul style="list-style-type: none"> • file • paper • papert • shortcut 	A list of objects that map the attributes or field names of your Dropbox files, Dropbox Paper, and shortcuts to Amazon Q index field names.
enableIdentityCrawler	Specify <code>true</code> to use the Amazon Q identity crawler to sync identity/principal information on users and groups with access to specific documents.
secretARN	<p>The Amazon Resource Name (ARN) of an AWS Secrets Manager secret that contains the key-value pairs required to connect to your Dropbox. The secret must contain a JSON structure with the following keys:</p> <pre style="border: 1px solid #ccc; border-radius: 10px; padding: 10px; background-color: #f9f9f9;">{ "appKey": "Dropbox app key", "appSecret": "Dropbox app secret", "accesstoken": "temporary access token or refresh token" }</pre>
additionalProperties	Additional configuration options for your content in your data source.
isCrawlAcl	Specify <code>true</code> to crawl access control information from documents.
fieldForUserId	Specify field to use for <code>UserId</code> for ACL crawling.

Configuration	Description
inclusionFileTypePatterns	A list of regular expression patterns to <i>include</i> specific file types in your Dropbox data source. Files that match the patterns are included in the index. Files that don't match the patterns are excluded from the index. If a file matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence and the file isn't included in the index.
exclusionFileTypePatterns	A list of regular expression patterns to <i>exclude</i> specific file types in your Dropbox data source. Files that match the patterns are excluded from the index. Files that don't match the patterns are included in the index. If a file matches both an exclusion and inclusion pattern, the exclusion pattern takes precedence and the file isn't included in the index.
exclusionFileNamePatterns	A list of regular expression patterns to <i>exclude</i> specific file names in your Dropbox data source. Files that match the patterns are excluded from the index. Files that don't match the patterns are included in the index. If a file matches both an exclusion and inclusion pattern, the exclusion pattern takes precedence and the file isn't included in the index.

Configuration	Description
exclusionFileNamePatterns	A list of regular expression patterns to <i>exclude</i> specific file names in your Dropbox data source. Files that match the patterns are excluded from the index. Files that don't match the patterns are included in the index. If a file matches both an exclusion and inclusion pattern, the exclusion pattern takes precedence and the file isn't included in the index.
<ul style="list-style-type: none"> • crawlFile • crawlPaper • crawlPapert • crawlShortcut 	true to index files in your Dropbox, Dropbox Paper documents, Dropbox Paper templates , and webpage shortcuts stored in your Dropbox.
type	The type of data source. Specify DROPBOX as your data source type.
useChangeLog	true to use the Dropbox change log to determine which documents require adding, updating, or deleting in the index. Depending on the change log's size, it may take longer for Amazon Q to use the change log than to scan all of your documents in your Dropbox.
tokenType	Specify your access token type: permanent or temporary access token. We recommend that you create a refresh access token that never expires in Dropbox rather than relying on a one-time access token that expires after 4 hours. You create an app and a refresh access token in the Dropbox developer console, and provide the access token in your secret.

Configuration	Description
version	The version of this template that's currently supported.

How Amazon Q Business connector crawls Dropbox ACLs

When you connect an Dropbox data source to Amazon Q Business, Amazon Q Business crawls ACL information attached to a document (user and group information) from your Dropbox instance. If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

The group and user IDs are mapped as follows:

- `_group_ids` – Group IDs exist in Dropbox on files where there are set access permissions. They're mapped from the names of the groups in Dropbox.
- `_user_id` – User IDs exist in Dropbox on files where there are set access permissions. They're mapped from the user emails as the IDs in Dropbox.

For more information, see:

- [Authorization](#)
- [Identity crawler](#)

Amazon Q Business Dropbox data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q Business enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a attribute mapping already available, or if you want to map additional document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q Dropbox connector supports the following entities and the associated reserved and custom attributes.

Supported entities and field mappings

- [Files](#)
- [Dropbox Paper](#)
- [Dropbox Paper Templates](#)
- [Shortcuts](#)

Files

Dropbox field name	Index field name	Description	Data type
sourceUrl	_source_uri	Default	String
category	_category	Default	String
fileName	dbx_file_name	Custom	String
fileId	dbx_id1	Custom	String
clientModifiedDate	dbx_client_modified	Custom	Date

Dropbox field name	Index field name	Description	Data type
serverModifiedDate	dbx_server_modified	Custom	Date
fileSize	dbx_file_size	Custom	Long (numeric)
pathDisplay	dbx_path_display	Custom	String
tags	dbx_tags	Custom	String

Dropbox Paper

Dropbox field name	Index field name	Description	Data type
sourceUrl	_source_uri	Default	String
category	_category	Default	String
fileName	dbx_file_name	Custom	String
fileId	dbx_id1	Custom	String
clientModifiedDate	dbx_client_modified	Custom	Date
serverModifiedDate	dbx_server_modified	Custom	Date
fileSize	dbx_file_size	Custom	Long (numeric)
pathDisplay	dbx_path_display	Custom	String
tags	dbx_tags	Custom	String

Dropbox Paper Templates

Dropbox field name	Index field name	Description	Data type
sourceUrl	_source_uri	Default	String
category	_category	Default	String

Dropbox field name	Index field name	Description	Data type
fileName	dbx_file_name	Custom	String
fileId	dbx_id1	Custom	String
clientModifiedDate	dbx_client_modified	Custom	Date
serverModifiedDate	dbx_server_modified	Custom	Date
fileSize	dbx_file_size	Custom	Long (numeric)
pathDisplay	dbx_path_display	Custom	String
tags	dbx_tags	Custom	String

Shortcuts

Dropbox field name	Index field name	Description	Data type
sourceUrl	_source_uri	Default	String
category	_category	Default	String
fileName	dbx_file_name	Custom	String
fileId	dbx_id1	Custom	String
clientModifiedDate	dbx_client_modified	Custom	Date
serverModifiedDate	dbx_server_modified	Custom	Date
fileSize	dbx_file_size	Custom	Long (numeric)
pathDisplay	dbx_path_display	Custom	String
tags	dbx_tags	Custom	String

IAM role for Amazon Q Business Dropbox connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the `BatchPutDocument` and `BatchDeleteDocument` operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToGetSecret",
      "Effect": "Allow",
      "Action": [
        "secretsmanager:GetSecretValue"
      ],
      "Resource": [
        "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
      ]
    },
    {
      "Sid": "AllowsAmazonQToDecryptSecret",
      "Effect": "Allow",
      "Action": [
        "kms:Decrypt"
      ],
      "Resource": [
```



```

    "arn:aws:kms:{{region}}:{{account_id}}:key/[key_id]"
  ],
  "Condition": {
    "StringLike": {
      "kms:ViaService": [
        "secretsmanager.*.amazonaws.com"
      ]
    }
  }
},
{
  "Sid": "AllowsAmazonQToIngestDocuments",
  "Effect": "Allow",
  "Action": [
    "qbusiness:BatchPutDocument",
    "qbusiness:BatchDeleteDocument"
  ],
  "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
},
{
  "Sid": "AllowsAmazonQToIngestPrincipalMapping",
  "Effect": "Allow",
  "Action": [
    "qbusiness:PutGroup",
    "qbusiness:CreateUser",
    "qbusiness>DeleteGroup",
    "qbusiness:UpdateUser",
    "qbusiness:ListGroups"
  ],
  "Resource": [
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}",
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}",
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}/data-source/*"
  ]
},
{
  "Sid": "AllowsAmazonQToCreateAndDeleteNI",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateNetworkInterface",
    "ec2>DeleteNetworkInterface"
  ]
}

```

```

    ],
    "Resource": [
      "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
      "arn:aws:ec2:{{region}}:{{account_id}}:security-group/[{{security_group}}]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2:DeleteNetworkInterface"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:RequestTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
      },
      "ForAllValues:StringEquals": {
        "aws:TagKeys": [
          "AMAZON_Q"
        ]
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateTags",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateTags"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringEquals": {
        "ec2:CreateAction": "CreateNetworkInterface"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterfacePermission"
    ],
  },

```

```

    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:ResourceTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
    "Effect": "Allow",
    "Action": [
      "ec2:DescribeNetworkInterfaces",
      "ec2:DescribeAvailabilityZones",
      "ec2:DescribeNetworkInterfaceAttribute",
      "ec2:DescribeVpcs",
      "ec2:DescribeRegions",
      "ec2:DescribeNetworkInterfacePermissions",
      "ec2:DescribeSubnets"
    ],
    "Resource": "*"
  }
]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "{{source_account}}"
        },
        "ArnEquals": {
          "aws:SourceArn": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/{{application_id}}"
        }
      }
    }
  ]
}

```

```
}  
  }  
    }  
  ]  
}
```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Connecting Drupal to Amazon Q Business

Drupal is an open-source content management system (CMS) that you can use to create websites and web applications. You can connect Drupal instance to Amazon Q Business—using either the AWS Management Console or the [CreateDataSource](#) API—and create an Amazon Q web experience.

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).
- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [Drupal connector overview](#)
- [Prerequisites for connecting Amazon Q Business to Drupal](#)
- [Connecting Amazon Q Business to Drupal using the console](#)
- [Connecting Amazon Q Business to Drupal using APIs](#)
- [How Amazon Q Business connector crawls Drupal ACLs](#)
- [Amazon Q Business Drupal data source connector field mappings](#)
- [IAM role for Amazon Q Business Drupal connector](#)
- [Known limitations for the Amazon Q Business Drupal connector](#)
- [Troubleshooting your Amazon Q Business Drupal connector](#)

Drupal connector overview

The following table gives an overview of the Amazon Q Business Drupal connector and its supported features.

Category	Feature	Support
Security	Authentication type	Basic, OAuth 2.0
	Authentication credentials	<p>Basic</p> <ul style="list-style-type: none"> • Username • Password • Client email • Private key <p>OAuth 2.0</p> <ul style="list-style-type: none"> • Username • Password • Client ID • Client Secret
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Identity crawling	Yes
	VPC	Yes
Crawl features	Custom metadata	Yes
	Entities	<p>Yes. The following entities are supported:</p> <ul style="list-style-type: none"> • Contents • Comments • Attachments

Category	Feature	Support
	Field mappings	Yes. Supports default and custom field mappings. For more information, see Field mappings .
	Filters	Yes. The following filters are supported: <ul style="list-style-type: none"> • Include/ exclude articles, article comments, and article attachments • Include/exclude basic pages, basic page comments, and basic page attachments • Include/exclude basic blocks, basic block comments, and basic block attachments • Include custom content types • Include custom blocks • Include/exclude content by article title, basic page title, basic block title, custom content title, custom block title, and file name
	Sync mode	Supports full and incremental sync.
	File types	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to Drupal

Before you begin, make sure that you have completed the following prerequisites.

In Drupal, make sure you have:

- Created a Drupal (Standard) Suite account and a user with an administrator role.
- Copied your Drupal site name and configured a host URL. For example, *https://<hostname>/<drupalsitename>*.
- Configured basic authentication credentials containing a user name (Drupal website login user name) and password (Drupal website password).

- **Recommended:** Configured an OAuth 2.0 credential token. Use this token along with your Drupal password grant, client id, client secret, user name (Drupal website login user name) and password (Drupal website password) to connect to Amazon Q.
- Added the following permissions in your Drupal account using an administrator role:
 - administer blocks
 - administer block_content display
 - administer block_content fields
 - administer block_content form display
 - administer views
 - view user email addresses
 - view own unpublished content
 - view page revisions
 - view article revisions
 - view all revisions
 - view the administration theme
 - access content
 - access content overview
 - access comments
 - search content
 - access files overview
 - access contextual links

 **Note**

If there are user defined content types or user defined block types, or any views and blocks are added to the Drupal website, they must be provided with administrator access.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.

- Stored your Drupal authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Connecting Amazon Q Business to Drupal using the console

The following procedure outlines how to connect Amazon Q Business to Drupal using the AWS Management Console.

Connecting Amazon Q to Drupal

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **Drupal** page, enter the following information:
6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. In **Source**, for **Host URL** – Enter the host URL of your Drupal site. For example, *https://<hostname>/<drupalstisename>*.
8. **SSL certificate location** – Enter the path to the SSL certificate stored in an Amazon S3 bucket. You use this to connect to Drupal with a secure SSL connection.

9. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

 **Note**

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

10. **Authentication** – Choose between **Basic authentication** and **OAuth 2.0 authentication** and then enter the following information for your **AWS Secrets Manager secret**.
 - a. **Basic authentication** – Enter the **User name**, (Drupal site user name), and **Password** (Drupal site password).
 - b. **OAuth 2.0 authentication** – Enter the **User name**, (Drupal site user name), and **Password** (Drupal site password).
11. **Configure VPC and security group – optional** – Choose whether you want to use a VPC. If you do, enter the following information:
 - a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
 - b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

12. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

13. In **Sync scope**, enter the following information:

- a. For **Select entities**:
 - **Articles**—Choose whether to crawl **Articles**, their comments **Comments**, and their **Attachments**.

- **Basic pages**—Choose whether to crawl **Basic pages**, their **Comments**, and their **Attachments**.
 - **Basic blocks**—Choose whether to crawl **Basic blocks**, their **Comments**, and their **Attachments**.
 - You can also choose to add and crawl **Custom content types** and **Custom Blocks**.
- b. (Optional) **Additional configuration** – Configure the following settings:
- **Regex pattern**—Add regular expression patterns to include or exclude specific entity titles and file names. You can add up to 100 patterns.
14. For **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.
- **Full sync**—Sync all content regardless of the previous sync status.
 - **New, modified, or deleted content sync**—Sync only new, modified, and deleted documents.
15. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
16. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
17. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:
- a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
 - b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

 **Note**

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

18. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

19. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

Note

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to Drupal using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the configuration parameter to provide a JSON schema with all other configuration information specific to your data source connector.

For an example of the API request, see [CreateDataSource](#) in the Amazon Q API Reference.

Drupal JSON schema

The following is the Drupal JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
      "type": "object",
      "properties": {
        "repositoryEndpointMetadata": {
          "type": "object",
          "properties": {
            "hostUrl": {
              "type": "string",
              "pattern": "https:.*"
            }
          }
        }
      }
    }
  }
}
```

```

    },
    "required": [
      "hostUrl"
    ]
  },
  "required": [
    "repositoryEndpointMetadata"
  ],
  "repositoryConfigurations": {
    "type": "object",
    "properties": {
      "content": {
        "type": "object",
        "properties": {
          "fieldMappings": {
            "type": "array",
            "items": [
              {
                "type": "object",
                "properties": {
                  "indexFieldName": {
                    "type": "string"
                  },
                  "indexFieldType": {
                    "type": "string",
                    "enum": [
                      "STRING",
                      "DATE"
                    ]
                  },
                  "dataSourceFieldName": {
                    "type": "string"
                  },
                  "dateFieldFormat": {
                    "type": "string",
                    "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
                  }
                }
              }
            ]
          },
          "required": [
            "indexFieldName",
            "indexFieldType",
            "dataSourceFieldName"
          ]
        }
      }
    }
  }
}

```

```
    ]
  }
]
},
"required": [
  "fieldMappings"
],
"comment": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": [
        {
          "type": "object",
          "properties": {
            "indexFieldName": {
              "type": "string"
            },
            "indexFieldType": {
              "type": "string",
              "enum": [
                "STRING",
                "DATE"
              ]
            },
            "dataSourceFieldName": {
              "type": "string"
            },
            "dateFieldFormat": {
              "type": "string",
              "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
            }
          }
        }
      ]
    },
    "required": [
      "indexFieldName",
      "indexFieldType",
      "dataSourceFieldName"
    ]
  }
}
```

```

    },
    "required": [
      "fieldMappings"
    ]
  },
  "attachment": {
    "type": "object",
    "properties": {
      "fieldMappings": {
        "type": "array",
        "items": [
          {
            "type": "object",
            "properties": {
              "indexFieldName": {
                "type": "string"
              },
              "indexFieldType": {
                "type": "string",
                "enum": [
                  "STRING",
                  "DATE"
                ]
              },
              "dataSourceFieldName": {
                "type": "string"
              },
              "dateFieldFormat": {
                "type": "string",
                "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
              }
            }
          }
        ]
      }
    }
  },
  "required": [
    "indexFieldName",
    "indexFieldType",
    "dataSourceFieldName"
  ]
}
},
"required": [
  "fieldMappings"
]

```

```
    }
  }
},
"additionalProperties": {
  "type": "object",
  "properties": {
    "isCrawlAcl": {
      "type": "boolean"
    },
    "fieldForUserId": {
      "type": "string"
    },
    "isCrawlArticle": {
      "type": "boolean"
    },
    "isCrawlBasicPage": {
      "type": "boolean"
    },
    "isCrawlBasicBlock": {
      "type": "boolean"
    },
    "crawlCustomContentTypesList": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    "crawlCustomBlockTypesList": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    "filePath": {
      "anyOf": [
        {
          "type": "string",
          "pattern": "s3:.*"
        },
        {
          "type": "string",
          "pattern": ""
        }
      ]
    }
  ]
}
```

```
},
"inclusionFileNamePatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"exclusionFileNamePatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"articleTitleInclusionPatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"articleTitleExclusionPatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"pageTitleInclusionPatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"pageTitleExclusionPatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"customContentTitleInclusionPatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"customContentTitleExclusionPatterns": {
```



```
"type": "array",
"items": {
  "type": "string"
}
},
"basicBlockTitleInclusionPatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"basicBlockTitleExclusionPatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"customBlockTitleInclusionPatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"customBlockTitleExclusionPatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"contentDefinitions": {
  "type": "array",
  "items": {
    "properties": {
      "contentType": {
        "type": "string"
      }
    },
    "fieldDefinition": {
      "type": "array",
      "items": [
        {
          "type": "object",
          "properties": {
            "machineName": {
              "type": "string"
            }
          }
        }
      ]
    }
  }
}
```

```
    },
    "type": {
      "type": "string"
    }
  },
  "required": [
    "machineName",
    "type"
  ]
}
]
},
"isCrawlComments": {
  "type": "boolean"
},
"isCrawlFiles": {
  "type": "boolean"
}
}
},
"required": [
  "contentType",
  "fieldDefinition",
  "isCrawlComments",
  "isCrawlFiles"
]
}
},
"required": []
},
"type": {
  "type": "string",
  "pattern": "DRUPAL"
},
"authType": {
  "type": "string",
  "enum": [
    "BASIC-AUTH",
    "OAUTH2"
  ]
},
"syncMode": {
  "type": "string",
  "enum": [
```

```

    "FORCED_FULL_CRAWL",
    "FULL_CRAWL",
    "CHANGE_LOG"
  ]
},
"enableIdentityCrawler": {
  "type": "boolean"
},
"secretArn": {
  "type": "string",
  "minLength": 20,
  "maxLength": 2048
}
},
"version": {
  "type": "string",
  "anyOf": [
    {
      "pattern": "1.0.0"
    }
  ]
},
"required": [
  "connectionConfiguration",
  "repositoryConfigurations",
  "syncMode",
  "additionalProperties",
  "secretArn",
  "type"
]
}

```

The following provides information on important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.
repositoryEndpointMetadata	The endpoint information for the data source.

Configuration	Description
hostUrl	The host URL of your Drupal website. For example, <i>https://<hostname>/<drupalsitename></i> .
repositoryConfigurations	Configuration information for the content of the data source.
<ul style="list-style-type: none"> content comment attachment 	A list of objects that map the attributes or field names of your Drupal files. The Drupal data source field names must exist in your Drupal custom metadata.
additionalProperties	Additional configuration options for your content in your data source.
isCrawlAcl	Specify true to crawl access control information from documents.
fieldForUserId	Specify field to use for UserId for ACL crawling.
<ul style="list-style-type: none"> inclusionFileNamePatterns articleTitleInclusionPatterns pageTitleInclusionPatterns customContentTitleInclusionPatterns basicBlockTitleInclusionPatterns customBlockTitleInclusionPatterns 	A list of regular expression patterns to <i>include</i> certain files in your Drupal data source. Files that match the patterns are included in the index. Files that don't match the patterns are excluded from the index. If a file matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence and the file isn't included in the index.

Configuration	Description
<ul style="list-style-type: none"> • exclusionFileNamePatterns • articleTitleExclusionPatterns • pageTitleExclusionPatterns • customContentTitleExclusionPatterns • basicBlockTitleExclusionPatterns • customBlockTitleExclusionPatterns 	<p>A list of regular expression patterns to <i>exclude</i> certain files in your Drupal data source. Files that match the patterns are excluded from the index. Files that don't match the patterns are included in the index. If a file matches both an exclusion and inclusion pattern, the exclusion pattern takes precedence and the file isn't included in the index.</p>
<p>contentDefinitions</p> <ul style="list-style-type: none"> • contentType • fieldDefinition • isCrawlComments • isCrawlFiles • isCrawlArticle • isCrawlBasicPage • isCrawlBasicBlock • isCrawlCustomContentTypesList 	<p>Specify the content types to crawl and whether to crawl comments and attachments for your selected content types.</p>
<p>type</p>	<p>The type of data source. Specify DRUPAL as your data source type.</p>
<p>authType</p>	<p>The type of authentication you are using, whether BASIC-AUTH or OAUTH2.</p>

Configuration	Description
syncMode	<p>Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose</p> <ul style="list-style-type: none">• <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index• <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index• <code>CHANGE_LOG</code> to incrementally crawl only new and modified content each time your data source syncs with your index
enableIdentityCrawler	<p><code>true</code> to activate identity crawler. Identity crawler is activated by default. Crawling identity information on users and groups with access to certain documents is useful for user context filtering. Search results are filtered based on the user or their group access to documents. See Identity crawler for more information.</p>

Configuration	Description
secretARN	<p>The Amazon Resource Name (ARN) of a Secrets Manager secret that contains the key-value pairs required to connect to your Drupal. The secret must contain a JSON structure with the following keys:</p> <p>If using basic authentication:</p> <pre data-bbox="829 569 1507 768"> { "user name": "user name", "passwords": "password" } </pre> <p>If using OAuth 2.0 authentication:</p> <pre data-bbox="829 877 1507 1150"> { "Client ID": "client_id", "Client secret": "client_secret", "user name": "user name", "password": "password" } </pre>
version	The version of this template that is currently supported.

How Amazon Q Business connector crawls Drupal ACLs

When you connect an Drupal data source to Amazon Q Business, Amazon Q Business crawls ACL information attached to a document (user and group information) from your Drupal instance. If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

Amazon Q gets the user and group information from the Drupal instance. The group and user IDs are mapped as follows:

- `_group_ids` – Group IDs exist in Drupal on files where there are set access permissions. They are mapped from the names of the groups in Drupal.

- `_user_id` – User IDs exist in Drupal on files where there are set access permissions. They are mapped from the user emails as the IDs in Drupal.

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business Drupal data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q Business enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a attribute mapping already available, or if you want to map additional document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q Drupal connector supports the following entities and the associated reserved and custom attributes.

Supported entities and field mappings

- [Contents](#)
- [Comments](#)
- [Attachments](#)

Contents

Drupal field name	Index field name	Entity	Category	Field type
title	dpl_title	All Entities	Default	String
sourceUrl	dpl_source_url	All Entities	Default	String
createdAt	dpl_created_date	All Entities	Default	Date
updatedAt	dpl_updated_date	All Entities	Default	Date
published	dpl_published	All Entities	Default	String
tag	dpl_tag	All Entities	Default	String
author	dpl_author	All Entities	Default	String
category	dpl_category	All Entities	Default	String
visibility	dpl_visibility	All Entities	Default	String
viewId	dpl_view_id	All Entities	Default	String

Comments

Drupal field name	Index field name	Entity	Category	Field type
title	dpl_comment_title	All Entities	Default	String
sourceUrl	dpl_source_url	All Entities	Default	String
createdAt	dpl_created_date	All Entities	Default	Date
updatedAt	dpl_updated_date	All Entities	Default	Date
approvedStatus	dpl_status	All Entities	Default	String
author	dpl_author	All Entities	Default	String
category	dpl_category	All Entities	Default	String
parentEntityId	dpl_parent_entity_id	All Entities	Default	String
visibility	dpl_visibility	All Entities	Default	String
viewId	dpl_view_id	All Entities	Default	String

Attachments

Drupal field name	Index field name	Entity	Category	Field type
fileName	dpl_file_name	All Entities	Default	String
sourceUrl	dpl_source_url	All Entities	Default	String

Drupal field name	Index field name	Entity	Category	Field type
createdAt	dpl_created_date	All Entities	Default	Date
updatedAt	dpl_updated_date	All Entities	Default	Date
status	dpl_status	All Entities	Default	String
fileType	dpl_file_type	All Entities	Default	String
fileSize	dpl_file_size	All Entities	Default	String
fileUploadedBy	dpl_file_uploaded_by	All Entities	Default	String
category	dpl_category	All Entities	Default	String
parentEntityId	dpl_parent_entity_id	All Entities	Default	String
visibility	dpl_visibility	All Entities	Default	String
viewId	dpl_view_id	All Entities	Default	String

IAM role for Amazon Q Business Drupal connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the BatchPutDocument and BatchDeleteDocument operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- Permission to access the SSL certificate stored in your Amazon S3 bucket.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AllowsAmazonQToGetS3Objects",
    "Action": [
      "s3:GetObject"
    ],
    "Resource": [
      "arn:aws:s3:::{{input_bucket_name}}/*"
    ],
    "Effect": "Allow",
    "Condition": {
      "StringEquals": {
        "aws:ResourceAccount": "{{account_id}}"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToGetSecret",
    "Effect": "Allow",
    "Action": [
      "secretsmanager:GetSecretValue"
    ],
    "Resource": [
      "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToDecryptSecret",
    "Effect": "Allow",
    "Action": [
```

```

        "kms:Decrypt"
    ],
    "Resource": [
        "arn:aws:kms:{{region}}:{{account_id}}:key/[key_id]"
    ],
    "Condition": {
        "StringLike": {
            "kms:ViaService": [
                "secretsmanager.*.amazonaws.com"
            ]
        }
    }
},
{
    "Sid": "AllowsAmazonQToIngestDocuments",
    "Effect": "Allow",
    "Action": [
        "qbusiness:BatchPutDocument",
        "qbusiness:BatchDeleteDocument"
    ],
    "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
},
{
    "Sid": "AllowsAmazonQToIngestPrincipalMapping",
    "Effect": "Allow",
    "Action": [
        "qbusiness:PutGroup",
        "qbusiness:CreateUser",
        "qbusiness>DeleteGroup",
        "qbusiness:UpdateUser",
        "qbusiness:ListGroup"
    ],
    "Resource": [
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}",
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}",
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}/data-source/*"
    ]
},
{
    "Sid": "AllowsAmazonQToCreateAndDeleteNI",

```

```

    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2>DeleteNetworkInterface"
    ],
    "Resource": [
      "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
      "arn:aws:ec2:{{region}}:{{account_id}}:security-group/
[[security_group]]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2>DeleteNetworkInterface"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:RequestTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
      },
      "ForAllValues:StringEquals": {
        "aws:TagKeys": [
          "AMAZON_Q"
        ]
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateTags",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateTags"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringEquals": {
        "ec2:CreateAction": "CreateNetworkInterface"
      }
    }
  }
},

```

```

    {
      "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
      "Effect": "Allow",
      "Action": [
        "ec2:CreateNetworkInterfacePermission"
      ],
      "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
      "Condition": {
        "StringLike": {
          "aws:ResourceTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
        }
      }
    },
    {
      "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
      "Effect": "Allow",
      "Action": [
        "ec2:DescribeNetworkInterfaces",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeNetworkInterfaceAttribute",
        "ec2:DescribeVpcs",
        "ec2:DescribeRegions",
        "ec2:DescribeNetworkInterfacePermissions",
        "ec2:DescribeSubnets"
      ],
      "Resource": "*"
    }
  ]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToAssumeRoleForServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
    }
  ]
}

```

```

    "Condition": {
      "StringEquals": {
        "aws:SourceAccount": "{{source_account}}"
      },
      "ArnLike": {
        "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
      }
    }
  }
]
}

```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Known limitations for the Amazon Q Business Drupal connector

- Drupal APIs have no official throttling limits.
- Java SDKs are not available for Drupal.
- Drupal data can be fetched only using native JSON API's.
- Content types not associated with any Drupal **View** cannot be crawled.
- You need administrator access to crawl data from Drupal **Blocks**.
- There is no JSON API available to create the user defined content type using HTTP verbs.
- The document body and comments for **Articles**, **Basic pages**, **Basic blocks**, user defined content type, and user defined block type, are displayed in HTML format. If the HTML content is not well-formed, then the HTML related tags will appear in the document body and comments and will be visible in Amazon Kendra search results.
- Content types and **Block** types without description or body will not be ingested into Amazon Q. Only **Comments** and **Attachments** of such **Content** or **Block** types will be ingested into your Amazon Q index.

Troubleshooting your Amazon Q Business Drupal connector

The following table provides information about error codes you may see for the Drupal connector and suggested troubleshooting actions.

Error code	Error message	Suggested resolution
DPL-5001	Invalid userName or password while validating credentials.	Provide valid values for the userName and password.
DPL-5002	Invalid userName or password while validating OAuth credentials.	Provide valid userName and password.
DPL-5003	Invalid clientId or clientSecret while validating OAuth credentials.	Provide valid clientId and clientSecret.
DPL-5100	Empty/null host URL	The hostUrl should not be null or empty.
DPL-5101	Null/empty user name.	Provide a valid userName.
DPL-5102	Null/empty password.	Provide a valid password.
DPL-5103	Null/empty ClientId.	Provide a valid clientId.
DPL-5104	Null/empty Client Secret.	Provide a valid clientSecret.
DPL-5105	Incorrect auth type in the repositoryAdditionalProperties.	The authType should be basicAuth or OAuth2.
DPL-5106	Null/empty auth type in the repositoryAdditionalProperties.	The authType should not be null or empty.
DPL-5107	Null/empty Repository Configurations.	The repositoryConfigurations should not be null or empty.
DPL-5108	Only String, String List, Date and Long formats are supported for the	Provide the supported format only for the indexFieldType in all the fieldMappings.

Error code	Error message	Suggested resolution
	indexFieldType in all the field mappings.	
DPL-5109	Not able to read the file from the S3 bucket location.	Check if valid JSON file is provided in the repositoryAdditionalProperties.
DPL-5110	Invalid Profile Name provided.	Provide a valid s3ProfileName in the repositoryAdditionalProperties. For example, default
DPL-5111	Null/empty Profile Name.	The s3ProfileName should not be null/empty in the repositoryAdditionalProperties.
DPL-5112	Invalid URI found during address validation.	Provide a valid hostUrl.
DPL-5131	Null/empty ContentTypes/BlockTypes in contentDefinitions.	Provide value for the contentType or blockType in contentDefinitions.
DPL-5132	Null/empty/unknown field definitions in the contentDefinitions.	The fieldDefinition should be an empty array only.
DPL-5133	Invalid field definitions in the contentDefinitions.	In the contentDefinitions, the fieldDefinition should be a json array only having machineName and type fields.
DPL-5134	Null/empty value found for the machineName in the fieldDefinition.	Provide value for the machineName in the fieldDefinition.
DPL-5135	Null/empty value found for the type in the fieldDefinition.	Provide value for the type in the fieldDefinition.

Error code	Error message	Suggested resolution
DPL-5136	Invalid isCrawlComments value.	isCrawlComments should be a boolean value true or false.
DPL-5137	Invalid isCrawlFiles value.	isCrawlFiles should be a boolean value true or false.
DPL-5138	The machineName is not found in the fieldDefinition.	Define the machineName as key in the fieldDefinition.
DPL-5139	The type is not found in the fieldDefinition.	Define the type as key in the fieldDefinition.
DPL-5151	Invalid inclusion file name patterns	Provide valid regex pattern in the inclusionFileNamePatterns.
DPL-5152	Invalid exclusion file name patterns.	Provide valid regex pattern in the exclusionFileNamePatterns.
DPL-5153	Invalid Article title inclusion patterns.	Provide valid regex pattern in the articleTitleInclusionPatterns.
DPL-5154	Invalid Article title exclusion patterns.	Please provide valid regex pattern in the articleTitleExclusionPatterns.
DPL-5155	Invalid Page title inclusion filter patterns.	Provide valid regex pattern in the pageTitleInclusionPatterns.
DPL-5156	Invalid Page title exclusion filter patterns.	Provide valid regex pattern in the pageTitleExclusionPatterns.
DPL-5157	Invalid Custom Content title inclusion filter patterns.	Provide valid regex pattern in the customContentTitleInclusionPatterns.

Error code	Error message	Suggested resolution
DPL-5158	Invalid Custom Content title exclusion filter patterns.	Provide valid regex pattern in the customContentTitleExclusionPatterns.
DPL-5159	Invalid Basic Block title inclusion filter patterns.	Provide valid regex pattern in the basicBlockTitleInclusionPatterns.
DPL-5160	Invalid Basic Block title exclusion filter patterns.	Provide valid regex pattern in the basicBlockTitleExclusionPatterns.
DPL-5161	Invalid Custom Block title inclusion filter patterns.	Provide valid regex pattern in the customBlockTitleInclusionPatterns.
DPL-5162	Invalid Custom Block title exclusion filter patterns.	Provide valid regex pattern in the customBlockTitleExclusionPatterns.
DPL-5200	IO Exception occurred while reading contents from Drupal.	Refer to the log for more details.
DPL-5201	Please try again later. Unknown exception occurred.	Unknown exception occurred. Refer to the log for more details.
DPL-5202	Issue occurred while initializing Views with acl info in the cache for content entity:	Issue with Views. Refer to the log for more details.
DPL-5203	Drupal Configuration found null during change access token of OAuth authentication.	Issue with OAuth Authentication. Refer to the log for more details.

Error code	Error message	Suggested resolution
DPL-5204	The generated access token is empty or null. Issue occurred while generating access token.	Access token should not be null/empty. Provide valid access token and try. If still issue exists, refer to the log for more details.
DPL-5205	User info with the given userName do not exist.	Verify the provided userName and correct it.
DPL-5206	The api response has empty data element.	Check the logs for details about empty response body.
DPL-5207	Either no records found or some issue with View filter criteria for content entity:	Refer to the log for more details.
DPL-5500	Drupal connection successful.	Drupal connection successful.

Connecting GitHub (Cloud) to Amazon Q Business

GitHub (Cloud) is a web-based hosting service for software development providing code storage and management services with version control. You can connect your GitHub (Cloud) instance to Amazon Q Business—using either the AWS Management Console, CLI, or the [CreateDataSource](#) API—and create an Amazon Q web experience.

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).
- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [GitHub \(Cloud\) connector overview](#)
- [Prerequisites for connecting Amazon Q Business to GitHub \(Cloud\)](#)
- [Connecting Amazon Q Business to GitHub \(Cloud\) using the console](#)
- [Connecting Amazon Q Business to GitHub \(Cloud\) using APIs](#)
- [How Amazon Q Business connector crawls GitHub \(Cloud\) ACLs](#)
- [Amazon Q Business GitHub \(Cloud\) data source connector field mappings](#)
- [IAM role for Amazon Q Business GitHub \(Cloud\) connector](#)

GitHub (Cloud) connector overview

The following table gives an overview of the Amazon Q Business GitHub (Cloud) connector and its supported features.

Category	Feature	Support
Security	Authentication type	OAuth 2.0
	Authentication credentials	<ul style="list-style-type: none"> • GitHub token
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Identity crawling	Yes
	VPC	Yes
Crawl features	Custom metadata	Yes
	Entities	<p>Yes. The following entities are supported:</p> <ul style="list-style-type: none"> • Repository • Repository Commit • Issue Document • Issue Comment • Issue Attachment • Pull Request Comment

Category	Feature	Support
		<ul style="list-style-type: none"> • Pull request Document • Pull Request Attachment
	Field mappings	Yes. Supports default and custom field mappings. For more information, see Field mappings .
	Filters	Yes. The following filters are supported: <ul style="list-style-type: none"> • Include select repositories • Include content by specific entities. • Include specific branched by name • Include/exclude content by file name, file type, and file path
	Sync mode	Supports full and incremental sync.
	File types	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to GitHub (Cloud)

Before you begin, make sure that you have completed the following prerequisites.

In GitHub (Cloud), make sure you have:

- Created a GitHub (Cloud) user with administrative permissions to the GitHub (Cloud) organization.
- Created a classic personal access token for authentication credentials. See [GitHub documentation on creating a personal access token](#).
- **Recommended:** Created an OAuth token for authentication credentials. Use OAuth token for better API throttle limits and connector performance. See [GitHub documentation on OAuth authorization](#).
- Noted the GitHub (Cloud) host URL for the type of GitHub (Cloud) service that you use. For example, the host URL for GitHub (Cloud) Cloud could be *https://api.github.com*.

- Noted the name of your organization for GitHub (Cloud) the GitHub Enterprise account you want to connect to. You can find your organization name by logging into GitHub (Cloud) desktop and selecting **Your organizations** under your profile picture dropdown.
- Added the following permissions in GitHub (Cloud):
 - repo:status
 - public_repo
 - repo:invite
 - read:org
 - user:email
 - read:user

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your GitHub (Cloud) authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Connecting Amazon Q Business to GitHub (Cloud) using the console

The following procedure outlines how to connect Amazon Q Business to GitHub (Cloud) using the AWS Management Console.

Connecting Amazon Q to GitHub (Cloud)

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.

2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **GitHub (Cloud)** page, enter the following information:
6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. **Source** – Choose your GitHub (Cloud) source details.
 - a. **GitHub (Cloud) source** – Choose GitHub (Cloud) Enterprise Cloud.
 - b. **GitHub (Cloud) host URL** – Enter the GitHub (Cloud) host name with the protocol ([http://](#) or [https://](#)). For example: *<https://api.github.com>*.
 - c. **GitHub (Cloud) organization name** – You can find your organization name when you log in to GitHub (Cloud) desktop and go to **Your organizations** under your profile picture dropdown.
8. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

 **Note**

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

9. **Authentication** – Enter the following information for your **AWS Secrets Manager secret**.
 - a. **Secret name** – A name for your secret.
 - b. **GitHub (Cloud) token** – Enter the access token you created in GitHub.
10. **Configure VPC and security group – optional** – Choose whether you want to use a VPC. If you do, enter the following information:
 - a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.

- b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

11. **Identity crawler** – Choose to activate Amazon Q identity crawler to sync identity information. For more information, see [Identity crawler](#).
12. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

13. In **Sync scope**, enter the following information:

- a. **Select repositories to crawl**—Select between crawling **All** repositories or **Select repositories**.

If you choose **Select repositories**, add names for the repositories in **Name of repository** and, optionally, the name of any specific branches in **Name of branch**.

- b. **Additional configuration – optional** – Configure the following settings:

- **Content types** – Select the file types you want to include.
- **Regex patterns** – Regular expression patterns to include or exclude certain files. You can add up to 100 patterns.

14. In **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.

- **Full sync** – Sync all content regardless of the previous sync status.
- **New or modified content sync** – Sync only new and modified documents.
- **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.

For more details, see [Sync mode](#).

15. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).

16. **Tags - *optional*** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
17. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:
 - a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
 - b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

 **Note**

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

18. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

19. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

 **Note**

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to GitHub (Cloud) using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the configuration parameter to provide a JSON schema with all other configuration information specific to your data source connector.

For an example of the API request, see [CreateDataSource](#) in the Amazon Q API Reference.

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the configuration parameter to provide a JSON schema with all other configuration information specific to your data source connector.

For an example of the API request, see [CreateDataSource](#) in the Amazon Q API Reference.

GitHub JSON schema

The following is the GitHub JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
      "type": "object",
      "properties": {
        "repositoryEndpointMetadata": {
          "type": "object",
          "properties": {
            "type": {
              "type": "string"
            },
            "hostUrl": {
              "type": "string",
              "pattern": "https://.*"
            },
            "organizationName": {
              "type": "string"
            }
          }
        },
        "required": [
          "type",
          "hostUrl",
          "organizationName"
        ]
      }
    }
  },
}
```

```

    "required": [
      "repositoryEndpointMetadata"
    ]
  },
  "repositoryConfigurations": {
    "type": "object",
    "properties": {
      "ghRepository": {
        "type": "object",
        "properties": {
          "fieldMappings": {
            "type": "array",
            "items": [
              {
                "type": "object",
                "properties": {
                  "indexFieldName": {
                    "type": "string"
                  },
                  "indexFieldType": {
                    "type": "string",
                    "enum": [
                      "STRING",
                      "STRING_LIST",
                      "DATE"
                    ]
                  },
                  "dataSourceFieldName": {
                    "type": "string"
                  },
                  "dateFieldFormat": {
                    "type": "string",
                    "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
                  }
                }
              }
            ]
          },
          "required": [
            "indexFieldName",
            "indexFieldType",
            "dataSourceFieldName"
          ]
        }
      }
    ]
  },
},

```

```

        "required": [
            "fieldMappings"
        ]
    },
    "ghCommit": {
        "type": "object",
        "properties": {
            "fieldMappings": {
                "type": "array",
                "items": [
                    {
                        "type": "object",
                        "properties": {
                            "indexFieldName": {
                                "type": "string"
                            },
                            "indexFieldType": {
                                "type": "string",
                                "enum": [
                                    "STRING",
                                    "STRING_LIST",
                                    "DATE"
                                ]
                            },
                            "dataSourceFieldName": {
                                "type": "string"
                            },
                            "dateFieldFormat": {
                                "type": "string",
                                "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
                            }
                        }
                    }
                ]
            },
            "required": [
                "indexFieldName",
                "indexFieldType",
                "dataSourceFieldName"
            ]
        }
    }
},
"required": [
    "fieldMappings"
]

```

```

    },
    "ghIssueDocument": {
      "type": "object",
      "properties": {
        "fieldMappings": {
          "type": "array",
          "items": [
            {
              "type": "object",
              "properties": {
                "indexFieldName": {
                  "type": "string"
                },
                "indexFieldType": {
                  "type": "string",
                  "enum": [
                    "STRING",
                    "STRING_LIST",
                    "DATE"
                  ]
                },
                "dataSourceFieldName": {
                  "type": "string"
                },
                "dateFieldFormat": {
                  "type": "string",
                  "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
                }
              }
            }
          ]
        },
        "required": [
          "indexFieldName",
          "indexFieldType",
          "dataSourceFieldName"
        ]
      }
    },
    "required": [
      "fieldMappings"
    ]
  },
  "ghIssueComment": {
    "type": "object",

```

```

    "properties": {
      "fieldMappings": {
        "type": "array",
        "items": [
          {
            "type": "object",
            "properties": {
              "indexFieldName": {
                "type": "string"
              },
              "indexFieldType": {
                "type": "string",
                "enum": [
                  "STRING",
                  "STRING_LIST",
                  "DATE"
                ]
              },
              "dataSourceFieldName": {
                "type": "string"
              },
              "dateFieldFormat": {
                "type": "string",
                "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
              }
            },
            "required": [
              "indexFieldName",
              "indexFieldType",
              "dataSourceFieldName"
            ]
          }
        ]
      },
      "required": [
        "fieldMappings"
      ]
    },
    "ghIssueAttachment": {
      "type": "object",
      "properties": {
        "fieldMappings": {
          "type": "array",

```



```

        "items": [
            {
                "type": "object",
                "properties": {
                    "indexFieldName": {
                        "type": "string"
                    },
                    "indexFieldType": {
                        "type": "string",
                        "enum": [
                            "STRING",
                            "STRING_LIST",
                            "DATE"
                        ]
                    },
                    "dataSourceFieldName": {
                        "type": "string"
                    },
                    "dateFieldFormat": {
                        "type": "string",
                        "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
                    }
                },
                "required": [
                    "indexFieldName",
                    "indexFieldType",
                    "dataSourceFieldName"
                ]
            }
        ]
    },
    "required": [
        "fieldMappings"
    ]
},
"ghPRDocument": {
    "type": "object",
    "properties": {
        "fieldMappings": {
            "type": "array",
            "items": [
                {
                    "type": "object",

```



```

    },
    "indexFieldType": {
      "type": "string",
      "enum": [
        "STRING",
        "STRING_LIST",
        "DATE"
      ]
    },
    "dataSourceFieldName": {
      "type": "string"
    },
    "dateFieldFormat": {
      "type": "string",
      "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
    }
  },
  "required": [
    "indexFieldName",
    "indexFieldType",
    "dataSourceFieldName"
  ]
}
]
}
},
"required": [
  "fieldMappings"
]
},
"ghPRAttachment": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": [
        {
          "type": "object",
          "properties": {
            "indexFieldName": {
              "type": "string"
            },
            "indexFieldType": {
              "type": "string",

```

```

        "enum": [
            "STRING",
            "STRING_LIST",
            "DATE"
        ]
    },
    "dataSourceFieldName": {
        "type": "string"
    },
    "dateFieldFormat": {
        "type": "string",
        "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
    }
},
"required": [
    "indexFieldName",
    "indexFieldType",
    "dataSourceFieldName"
]
}
}
},
"required": [
    "fieldMappings"
]
}
},
"additionalProperties": {
    "type": "object",
    "properties": {
        "isCrawlAcl": {
            "type": "boolean"
        },
        "fieldForUserId": {
            "type": "string"
        },
        "crawlRepository": {
            "type": "boolean"
        },
        "crawlRepositoryDocuments": {
            "type": "boolean"
        }
    }
},

```

```
"crawlIssue": {
  "type": "boolean"
},
"crawlIssueComment": {
  "type": "boolean"
},
"crawlIssueCommentAttachment": {
  "type": "boolean"
},
"crawlPullRequest": {
  "type": "boolean"
},
"crawlPullRequestComment": {
  "type": "boolean"
},
"crawlPullRequestCommentAttachment": {
  "type": "boolean"
},
"repositoryFilter": {
  "type": "array",
  "items": [
    {
      "type": "object",
      "properties": {
        "repositoryName": {
          "type": "string"
        },
        "branchNameList": {
          "type": "array",
          "items": {
            "type": "string"
          }
        }
      }
    }
  ]
},
"inclusionFolderNamePatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"inclusionFileTypePatterns": {
```

```

        "type": "array",
        "items": {
            "type": "string"
        }
    },
    "inclusionFileNamePatterns": {
        "type": "array",
        "items": {
            "type": "string"
        }
    },
    "exclusionFolderNamePatterns": {
        "type": "array",
        "items": {
            "type": "string"
        }
    },
    "exclusionFileTypePatterns": {
        "type": "array",
        "items": {
            "type": "string"
        }
    },
    "exclusionFileNamePatterns": {
        "type": "array",
        "items": {
            "type": "string"
        }
    }
},
"required": []
},
"type": {
    "type": "string",
    "pattern": "GITHUB"
},
"syncMode": {
    "type": "string",
    "enum": [
        "FULL_CRAWL",
        "FORCED_FULL_CRAWL",
        "CHANGE_LOG"
    ]
},
},

```

```

    "enableIdentityCrawler": {
      "type": "boolean"
    },
    "secretArn": {
      "type": "string",
      "minLength": 20,
      "maxLength": 2048
    }
  },
  "version": {
    "type": "string",
    "anyOf": [
      {
        "pattern": "1.0.0"
      }
    ]
  },
  "required": [
    "connectionConfiguration",
    "repositoryConfigurations",
    "syncMode",
    "additionalProperties",
    "enableIdentityCrawler"
  ]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.
repositoryEndpointMetadata	The endpoint information for the data source.
hostUrl	The GitHub (Cloud) host URL. For example, if you use GitHub (Cloud) Enterprise Cloud: <i>https://api.github.com</i> . Or, if you use GitHub (Cloud) Enterprise Server: <i>https://on-prem-host-url/api/v3/</i> .

Configuration	Description
organizationName	You can find your organization name when you log in to GitHub (Cloud) desktop and go to Your organizations under your profile picture dropdown.
repositoryConfigurations	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings.
<ul style="list-style-type: none"> • ghRepository • ghCommit • ghIssueDocument • ghIssueComment • ghIssueAttachment • ghPRDocument • ghPRComment • ghPRAttachment 	A list of objects that map the attributes or field names of your GitHub pages and assets to Amazon Q index field names.
additionalProperties	Additional configuration options for your content in your data source.
isCrawlAcl	Specify <code>true</code> to crawl access control information from documents.
fieldForUserId	Specify field to use for <code>UserId</code> for ACL crawling.
repositoryFilter	A list of names of the specific repositories and branch names you want to index.
crawlRepository	Specify <code>true</code> to crawl repositories.
crawlRepositoryDocuments	Specify <code>true</code> to crawl repository documents.
crawlIssue	Specify <code>true</code> to crawl issues.

Configuration	Description
<code>crawlIssueComment</code>	Specify <code>true</code> to crawl issue comments.
<code>crawlIssueCommentAttachment</code>	Specify <code>true</code> to crawl issue comment attachments.
<code>crawlPullRequest</code>	Specify <code>true</code> to crawl pull requests.
<code>crawlPullRequestComment</code>	Specify <code>true</code> to crawl pull request comments.
<code>crawlPullRequestCommentAttachment</code>	Specify <code>true</code> to crawl pull request comment attachments.
<ul style="list-style-type: none"> <code>inclusionFolderNamePatterns</code> <code>inclusionFileTypePatterns</code> <code>inclusionFileNamePatterns</code> 	A list of regular expression patterns to include specific content in your GitHub data source. Content that matches the patterns are included in the index. Content that doesn't match the patterns are excluded from the index. If any content matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence, and the content isn't included in the index.
<ul style="list-style-type: none"> <code>exclusionFolderNamePatterns</code> <code>exclusionFileTypePatterns</code> <code>exclusionFileNamePatterns</code> 	A list of regular expression patterns to exclude specific content in your GitHub data source. Content that matches the patterns are included in the index. Content that doesn't match the patterns are excluded from the index. If any content matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence, and the content isn't included in the index.
<code>type</code>	The type of data source. Specify <code>GITHUB</code> as your data source type.

Configuration	Description
<code>enableIdentityCrawler</code>	Specify <code>true</code> to use the Amazon Q identity crawler to sync identity/principal information on users and groups with access to specific documents.
<code>syncMode</code>	<p>Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose between the following options:</p> <ul style="list-style-type: none">• Use <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index.• Use <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index.• Use <code>CHANGE_LOG</code> to incrementally crawl only new and modified content each time your data source syncs with your index.
<code>secretArn</code>	<p>The Amazon Resource Name (ARN) of an AWS Secrets Manager secret that contains the key-value pairs required to connect to your GitHub (Cloud). The secret must contain a JSON structure with the following keys:</p> <pre data-bbox="829 1514 1507 1675">{ "personalToken": " <i>token</i>" }</pre>
<code>version</code>	The version of this template that's currently supported.

How Amazon Q Business connector crawls GitHub (Cloud) ACLs

When you connect an GitHub (Cloud) data source to Amazon Q Business, Amazon Q crawls ACL information attached to a document (user and group information) from your GitHub (Cloud) instance. If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

The user IDs are mapped as follows:

- `_user_id` – User IDs exist in GitHub on files where there are set access permissions. They are mapped from the user emails as the IDs in GitHub.

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business GitHub (Cloud) data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q Business enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a attribute mapping already available, or if you want to map additional document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q GitHub connector supports the following entities and the associated reserved and custom attributes.

Note

You can map any GitHub field to the document title or document body Amazon Q reserved/default index fields.

Supported entities and field mappings

- [Repository](#)
- [Repository Commit](#)
- [Issue Document](#)
- [Issue Comment](#)
- [Issue Attachment](#)
- [Pull Request Comment](#)
- [Pull Request Document](#)
- [Pull Request Attachment](#)

Repository

GitHub field name	Index field name	Description	Data type
Description	_document_body	Default	String
repositoryName	gh_repository_name	Custom	String

GitHub field name	Index field name	Description	Data type
repositoryVisibility	gh_repository_visibility	Custom	String
category	_category	Default	String
owner	_authors	Default	String list
sourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date

Repository Commit

GitHub field name	Index field name	Description	Data type
Description	_document_body	Default	String
repositoryName	gh_repository_name	Custom	String
repositoryVisibility	gh_repository_visibility	Custom	String
category	_category	Default	String
fileType	_file_type	Default	String
owner	_authors	Default	String list
sourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date
fileName	gh_file_name	Default	String

GitHub field name	Index field name	Description	Data type
fileSize	gh_size	Default	Long (numeric)
branchName	gh_branch_name	Default	String

Issue Document

GitHub field name	Index field name	Description	Data type
repositoryName	gh_repository_name	Custom	String
repositoryVisibility	gh_repository_visibility	Custom	String
category	_category	Default	String
issueNumber	gh_issue_number	Custom	Long (numeric)
issueTitle	gh_issue_title	Custom	String
owner	_authors	Default	String list
fileType	_file_type	Default	String
issueSourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date
issueFileName	gh_file_name	Custom	String
issueState	gh_issue_state	Custom	String
issueLabel	gh_issue_labels	Default	String list
issueAssignee	gh_issue_assignee	Default	String list

Issue Comment

GitHub field name	Index field name	Description	Data type
repositoryName	gh_repository_name	Custom	String
repositoryVisibility	gh_repository_visibility	Custom	String
category	_category	Default	String
issueNumber	gh_issue_number	Custom	Long (numeric)
issueTitle	gh_issue_title	Custom	String
owner	_authors	Default	String list
issueSourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date
issueState	gh_issue_state	Custom	String
issueLabel	gh_issue_labels	Default	String list
issueAssignee	gh_issue_assignee	Default	String list

Issue Attachment

GitHub field name	Index field name	Description	Data type
repositoryName	gh_repository_name	Custom	String
repositoryVisibility	gh_repository_visibility	Custom	String
category	_category	Default	String

GitHub field name	Index field name	Description	Data type
issueNumber	gh_issue_number	Custom	Long (numeric)
issueTitle	gh_issue_title	Custom	String
owner	_authors	Default	String list
issueSourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date
issueFileName	gh_file_name	Custom	String
issueFileType	_file_type	Custom	String
issueState	gh_issue_state	Custom	String
issueLabel	gh_issue_labels	Default	String list
issueAssignee	gh_issue_assignee	Default	String list

Pull Request Comment

GitHub field name	Index field name	Description	Data type
repositoryName	gh_repository_name	Custom	String
repositoryVisibility	gh_repository_visibility	Custom	String
category	_category	Default	String
PRNumber	gh_pr_number	Custom	Long (numeric)
PRTitle	gh_pr_title	Custom	String
owner	_authors	Default	String list

GitHub field name	Index field name	Description	Data type
PRSourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date
PRState	gh_pr_state	Custom	String
PRLabel	gh_pr_labels	Default	String list
PRAssignee	gh_pr_assignee	Default	String list

Pull Request Document

GitHub field name	Index field name	Description	Data type
repositoryName	gh_repository_name	Custom	String
repositoryVisibility	gh_repository_visibility	Custom	String
category	_category	Default	String
PRNumber	gh_number	Custom	Long (numeric)
PRTitle	gh_pr_title	Custom	String
owner	_authors	Default	String list
PRSourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date
PRFileName	gh_file_name	Custom	String
PRFileType	_file_type	Custom	String

GitHub field name	Index field name	Description	Data type
PRState	gh_pr_state	Custom	String
PRLabel	gh_pr_labels	Default	String list
PRAssignee	gh_pr_assignee	Default	String list

Pull Request Attachment

GitHub field name	Index field name	Description	Data type
repositoryName	gh_repository_name	Custom	String
repositoryVisibility	gh_repository_visibility	Custom	String
category	_category	Default	String
PRNumber	gh_number	Custom	Long (numeric)
PRTitle	gh_pr_title	Custom	String
owner	_authors	Default	String list
PRSourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date
PRFileName	gh_file_name	Custom	String
PRFileType	_file_type	Custom	String
PRState	gh_pr_state	Custom	String
PRLabel	gh_pr_labels	Default	String list
PRAssignee	gh_pr_assignee	Default	String list

IAM role for Amazon Q Business GitHub (Cloud) connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the `BatchPutDocument` and `BatchDeleteDocument` operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToGetSecret",
      "Effect": "Allow",
      "Action": [
        "secretsmanager:GetSecretValue"
      ],
      "Resource": [
        "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
      ]
    },
    {
      "Sid": "AllowsAmazonQToDecryptSecret",
      "Effect": "Allow",
      "Action": [
        "kms:Decrypt"
      ],
      "Resource": [
```

```

    "arn:aws:kms:{{region}}:{{account_id}}:key/[{{key_id}}]"
  ],
  "Condition": {
    "StringLike": {
      "kms:ViaService": [
        "secretsmanager.*.amazonaws.com"
      ]
    }
  }
},
{
  "Sid": "AllowsAmazonQToIngestDocuments",
  "Effect": "Allow",
  "Action": [
    "qbusiness:BatchPutDocument",
    "qbusiness:BatchDeleteDocument"
  ],
  "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
},
{
  "Sid": "AllowsAmazonQToIngestPrincipalMapping",
  "Effect": "Allow",
  "Action": [
    "qbusiness:PutGroup",
    "qbusiness:CreateUser",
    "qbusiness>DeleteGroup",
    "qbusiness:UpdateUser",
    "qbusiness:ListGroup"
  ],
  "Resource": [
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}",
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}",
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}/data-source/*"
  ]
},
{
  "Sid": "AllowsAmazonQToCreateAndDeleteNI",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateNetworkInterface",
    "ec2>DeleteNetworkInterface"
  ]
}

```

```

    ],
    "Resource": [
      "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
      "arn:aws:ec2:{{region}}:{{account_id}}:security-group/[{{security_group}}]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2:DeleteNetworkInterface"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:RequestTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
      },
      "ForAllValues:StringEquals": {
        "aws:TagKeys": [
          "AMAZON_Q"
        ]
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateTags",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateTags"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringEquals": {
        "ec2:CreateAction": "CreateNetworkInterface"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterfacePermission"
    ],
  },

```

```

    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:ResourceTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
    "Effect": "Allow",
    "Action": [
      "ec2:DescribeNetworkInterfaces",
      "ec2:DescribeAvailabilityZones",
      "ec2:DescribeNetworkInterfaceAttribute",
      "ec2:DescribeVpcs",
      "ec2:DescribeRegions",
      "ec2:DescribeNetworkInterfacePermissions",
      "ec2:DescribeSubnets"
    ],
    "Resource": "*"
  }
]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "{{source_account}}"
        },
        "ArnEquals": {
          "aws:SourceArn": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/{{application_id}}"
        }
      }
    }
  ]
}

```

```
}
  }
    }
  ]
}
```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Connecting GitHub (Server) to Amazon Q Business

GitHub (Server) is a web-based hosting service for software development providing code storage and management services with version control. You can connect your GitHub (Server) instance to Amazon Q Business—using either the AWS Management Console, CLI, or the [CreateDataSource](#) API—and create an Amazon Q web experience.

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).
- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [GitHub \(Server\) connector overview](#)
- [Prerequisites for connecting Amazon Q Business to GitHub \(Server\)](#)
- [Connecting Amazon Q Business to GitHub \(Server\) using the console](#)
- [Connecting Amazon Q Business to GitHub \(Server\) using APIs](#)
- [How Amazon Q Business connector crawls GitHub \(Server\) ACLs](#)
- [Amazon Q Business GitHub \(Server\) data source connector field mappings](#)
- [IAM role for Amazon Q Business GitHub \(Server\) connector](#)

GitHub (Server) connector overview

The following table gives an overview of the Amazon Q Business GitHub (Server) connector and its supported features.

Category	Feature	Support
Security	Authentication type	OAuth 2.0
	Authentication credentials	<ul style="list-style-type: none"> • GitHub token
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Identity crawling	Yes
	VPC	Yes
Crawl features	Custom metadata	Yes
	Entities	<p>Yes. The following entities are supported:</p> <ul style="list-style-type: none"> • Repository • Repository Commit • Issue Document • Issue Comment • Issue Attachment • Pull Request Comment • Pull request Document • Pull Request Attachment
	Field mappings	Yes. Supports default and custom field mappings. For more information, see Field mappings .
	Filters	<p>Yes. The following filters are supported:</p> <ul style="list-style-type: none"> • Include select repositories • Include content by specific entities. • Include specific branched by name • Include/exclude content by file name, file type, and file path

Category	Feature	Support
	Sync mode	Supports full and incremental sync.
	File types	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to GitHub (Server)

Before you begin, make sure that you have completed the following prerequisites.

In GitHub (Server), make sure you have:

- Created a GitHub (Server) user with administrative permissions to the GitHub (Server) organization.
- Created a classic personal access token for authentication credentials. See [GitHub documentation on creating a personal access token](#).
- **Recommended:** Created an OAuth token for authentication credentials. Use OAuth token for better API throttle limits and connector performance. See [GitHub documentation on OAuth authorization](#).
- Noted the GitHub (Server) host URL for the type of GitHub (Server) service that you use. For example, the host URL for GitHub (Server) Server could be *https://on-prem-host-url/api/v3/*.
- Noted the name of your organization for GitHub (Server) the GitHub Enterprise account you want to connect to. You can find your organization name by logging into GitHub (Server) desktop and selecting **Your organizations** under your profile picture dropdown.
- Added the following permissions in GitHub (Server):
 - repo:status
 - public_repo
 - repo:invite
 - read:org
 - user:email
 - read:user

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your GitHub (Server) authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

 **Note**

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Connecting Amazon Q Business to GitHub (Server) using the console

The following procedure outlines how to connect Amazon Q Business to GitHub (Server) using the AWS Management Console.

Connecting Amazon Q to GitHub (Server)

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **GitHub (Server)** page, enter the following information:
6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. **Source** – Choose your GitHub (Server) source details.
 - a. **GitHub (Server) source** – Choose GitHub (Server) Enterprise Cloud.
 - b. **GitHub (Server) host URL** – Enter the GitHub (Server) host name with the protocol (http:// or https://). For example: *https://on-prem-host-url/api/v3/*.

- c. **GitHub (Server) organization name** – You can find your organization name when you log in to GitHub (Server) desktop and go to **Your organizations** under your profile picture dropdown.
 - d. **SSL certificate location**— Enter the path to the SSL certificate stored in an Amazon S3 bucket. You use this to connect to Github (Server) with a secure SSL connection.
8. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

 **Note**

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

9. **Authentication** – Enter the following information for your **AWS Secrets Manager secret**.
- a. **Secret name** – A name for your secret.
 - b. **GitHub (Server) token** – Enter the access token you created in GitHub.
10. **Configure VPC and security group – optional** – Choose whether you want to use a VPC. If you do, enter the following information:
- a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
 - b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

11. **Identity crawler** – Choose to activate Amazon Q identity crawler to sync identity information. For more information, see [Identity crawler](#).
12. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

13. In **Sync scope**, enter the following information:

- a. **Select repositories to crawl**—Select between crawling **All** repositories or **Select repositories**.

If you choose **Select repositories**, add names for the repositories in **Name of repository** and, optionally, the name of any specific branches in **Name of branch**.

- b. **Additional configuration – optional** – Configure the following settings:
 - **Content types** – Select the file types you want to include.
 - **Regex patterns** – Regular expression patterns to include or exclude certain files. You can add up to 100 patterns.

14. In **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.

- **Full sync** – Sync all content regardless of the previous sync status.
- **New or modified content sync** – Sync only new and modified documents.
- **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.

For more details, see [Sync mode](#).

15. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).

16. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.

17. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:

- a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
- b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

Note

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

18. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

19. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

Note

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to GitHub (Server) using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the `configuration` parameter to provide a JSON schema with all other configuration information specific to your data source connector.

For an example of the API request, see [CreateDataSource](#) in the Amazon Q API Reference.

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the `configuration` parameter to provide a JSON schema with all other configuration information specific to your data source connector.

For an example of the API request, see [CreateDataSource](#) in the Amazon Q API Reference.

GitHub JSON schema

The following is the GitHub JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
      "type": "object",
      "properties": {
        "repositoryEndpointMetadata": {
          "type": "object",
          "properties": {
            "type": {
              "type": "string"
            },
            "hostUrl": {
              "type": "string",
              "pattern": "https://.*"
            },
            "organizationName": {
              "type": "string"
            }
          }
        },
        "required": [
          "type",
          "hostUrl",
          "organizationName"
        ]
      }
    },
    "required": [
      "repositoryEndpointMetadata"
    ]
  },
  "repositoryConfigurations": {
    "type": "object",
    "properties": {
      "ghRepository": {
        "type": "object",
        "properties": {
          "fieldMappings": {
            "type": "array",

```

```

        "items": [
            {
                "type": "object",
                "properties": {
                    "indexFieldName": {
                        "type": "string"
                    },
                    "indexFieldType": {
                        "type": "string",
                        "enum": [
                            "STRING",
                            "STRING_LIST",
                            "DATE"
                        ]
                    },
                    "dataSourceFieldName": {
                        "type": "string"
                    },
                    "dateFieldFormat": {
                        "type": "string",
                        "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
                    }
                },
                "required": [
                    "indexFieldName",
                    "indexFieldType",
                    "dataSourceFieldName"
                ]
            }
        ]
    },
    "required": [
        "fieldMappings"
    ]
},
"ghCommit": {
    "type": "object",
    "properties": {
        "fieldMappings": {
            "type": "array",
            "items": [
                {
                    "type": "object",

```



```

        },
        "indexFieldType": {
            "type": "string",
            "enum": [
                "STRING",
                "STRING_LIST",
                "DATE"
            ]
        },
        "dataSourceFieldName": {
            "type": "string"
        },
        "dateFieldFormat": {
            "type": "string",
            "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
        }
    },
    "required": [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
    ]
}
]
}
},
"required": [
    "fieldMappings"
]
},
"ghIssueComment": {
    "type": "object",
    "properties": {
        "fieldMappings": {
            "type": "array",
            "items": [
                {
                    "type": "object",
                    "properties": {
                        "indexFieldName": {
                            "type": "string"
                        },
                    },
                    "indexFieldType": {
                        "type": "string",

```

```

        "enum": [
            "STRING",
            "STRING_LIST",
            "DATE"
        ]
    },
    "dataSourceFieldName": {
        "type": "string"
    },
    "dateFieldFormat": {
        "type": "string",
        "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
    }
},
"required": [
    "indexFieldName",
    "indexFieldType",
    "dataSourceFieldName"
]
}
]
}
},
"required": [
    "fieldMappings"
]
},
"ghIssueAttachment": {
    "type": "object",
    "properties": {
        "fieldMappings": {
            "type": "array",
            "items": [
                {
                    "type": "object",
                    "properties": {
                        "indexFieldName": {
                            "type": "string"
                        },
                    },
                    "indexFieldType": {
                        "type": "string",
                        "enum": [
                            "STRING",
                            "STRING_LIST",

```

```

        "DATE"
      ]
    },
    "dataSourceFieldName": {
      "type": "string"
    },
    "dateFieldFormat": {
      "type": "string",
      "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
    }
  },
  "required": [
    "indexFieldName",
    "indexFieldType",
    "dataSourceFieldName"
  ]
}
]
}
},
"required": [
  "fieldMappings"
]
},
"ghPRDocument": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": [
        {
          "type": "object",
          "properties": {
            "indexFieldName": {
              "type": "string"
            },
            "indexFieldType": {
              "type": "string",
              "enum": [
                "STRING",
                "STRING_LIST",
                "DATE"
              ]
            }
          }
        }
      ]
    }
  }
},

```

```

        "dataSourceFieldName": {
            "type": "string"
        },
        "dateFieldFormat": {
            "type": "string",
            "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
        }
    },
    "required": [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
    ]
}
]
}
},
"required": [
    "fieldMappings"
]
},
"ghPRComment": {
    "type": "object",
    "properties": {
        "fieldMappings": {
            "type": "array",
            "items": [
                {
                    "type": "object",
                    "properties": {
                        "indexFieldName": {
                            "type": "string"
                        },
                        "indexFieldType": {
                            "type": "string",
                            "enum": [
                                "STRING",
                                "STRING_LIST",
                                "DATE"
                            ]
                        }
                    },
                    "dataSourceFieldName": {
                        "type": "string"
                    }
                }
            ]
        }
    }
}

```

```

        "dateFieldFormat": {
            "type": "string",
            "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
        }
    },
    "required": [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
    ]
}
]
}
},
"required": [
    "fieldMappings"
]
},
"ghPRAttachment": {
    "type": "object",
    "properties": {
        "fieldMappings": {
            "type": "array",
            "items": [
                {
                    "type": "object",
                    "properties": {
                        "indexFieldName": {
                            "type": "string"
                        },
                    },
                    "indexFieldType": {
                        "type": "string",
                        "enum": [
                            "STRING",
                            "STRING_LIST",
                            "DATE"
                        ]
                    },
                }
            ],
        },
        "dataSourceFieldName": {
            "type": "string"
        },
        "dateFieldFormat": {
            "type": "string",
            "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
        }
    }
}

```

```
        }
      },
      "required": [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
      ]
    }
  ]
},
"required": [
  "fieldMappings"
]
}
},
"additionalProperties": {
  "type": "object",
  "properties": {
    "isCrawlAcl": {
      "type": "boolean"
    },
    "fieldForUserId": {
      "type": "string"
    },
    "crawlRepository": {
      "type": "boolean"
    },
    "crawlRepositoryDocuments": {
      "type": "boolean"
    },
    "crawlIssue": {
      "type": "boolean"
    },
    "crawlIssueComment": {
      "type": "boolean"
    },
    "crawlIssueCommentAttachment": {
      "type": "boolean"
    },
    "crawlPullRequest": {
      "type": "boolean"
    }
  },
}
```

```
"crawlPullRequestComment": {
  "type": "boolean"
},
"crawlPullRequestCommentAttachment": {
  "type": "boolean"
},
"repositoryFilter": {
  "type": "array",
  "items": [
    {
      "type": "object",
      "properties": {
        "repositoryName": {
          "type": "string"
        },
        "branchNameList": {
          "type": "array",
          "items": {
            "type": "string"
          }
        }
      }
    }
  ]
},
"inclusionFolderNamePatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"inclusionFileTypePatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"inclusionFileNamePatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"exclusionFolderNamePatterns": {
```

```
        "type": "array",
        "items": {
            "type": "string"
        }
    },
    "exclusionFileTypePatterns": {
        "type": "array",
        "items": {
            "type": "string"
        }
    },
    "exclusionFileNamePatterns": {
        "type": "array",
        "items": {
            "type": "string"
        }
    }
},
"required": []
},
"type": {
    "type": "string",
    "pattern": "GITHUB"
},
"syncMode": {
    "type": "string",
    "enum": [
        "FULL_CRAWL",
        "FORCED_FULL_CRAWL",
        "CHANGE_LOG"
    ]
},
"enableIdentityCrawler": {
    "type": "boolean"
},
"secretArn": {
    "type": "string",
    "minLength": 20,
    "maxLength": 2048
}
},
"version": {
    "type": "string",
    "anyOf": [
```



```

        {
            "pattern": "1.0.0"
        }
    ],
    "required": [
        "connectionConfiguration",
        "repositoryConfigurations",
        "syncMode",
        "additionalProperties",
        "enableIdentityCrawler"
    ]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.
repositoryEndpointMetadata	The endpoint information for the data source.
hostUrl	The GitHub (Server) host URL. For example, if you use GitHub (Server) Enterprise Cloud: <i>https://api.github.com</i> . Or, if you use GitHub (Server) Enterprise Server: <i>https://on-prem-host-url/api/v3/</i> .
organizationName	You can find your organization name when you log in to GitHub (Server) desktop and go to Your organizations under your profile picture dropdown.
repositoryConfigurations	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings.

Configuration	Description
<ul style="list-style-type: none"> ghRepository ghCommit ghIssueDocument ghIssueComment ghIssueAttachment ghPRDocument ghPRComment ghPRAttachment 	A list of objects that map the attributes or field names of your GitHub pages and assets to Amazon Q index field names.
additionalProperties	Additional configuration options for your content in your data source.
isCrawlAcl	Specify true to crawl access control information from documents.
fieldForUserId	Specify field to use for UserId for ACL crawling.
repositoryFilter	A list of names of the specific repositories and branch names you want to index.
crawlRepository	Specify true to crawl repositories.
crawlRepositoryDocuments	Specify true to crawl repository documents.
crawlIssue	Specify true to crawl issues.
crawlIssueComment	Specify true to crawl issue comments.
crawlIssueCommentAttachment	Specify true to crawl issue comment attachments.
crawlPullRequest	Specify true to crawl pull requests.
crawlPullRequestComment	Specify true to crawl pull request comments.

Configuration	Description
<code>crawlPullRequestCommentAttachment</code>	Specify <code>true</code> to crawl pull request comment attachments.
<ul style="list-style-type: none"> • <code>inclusionFolderNamePatterns</code> • <code>inclusionFileTypePatterns</code> • <code>inclusionFileNamePatterns</code> 	A list of regular expression patterns to include specific content in your GitHub data source. Content that matches the patterns are included in the index. Content that doesn't match the patterns are excluded from the index. If any content matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence, and the content isn't included in the index.
<ul style="list-style-type: none"> • <code>exclusionFolderNamePatterns</code> • <code>exclusionFileTypePatterns</code> • <code>exclusionFileNamePatterns</code> 	A list of regular expression patterns to exclude specific content in your GitHub data source. Content that matches the patterns are included in the index. Content that doesn't match the patterns are excluded from the index. If any content matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence, and the content isn't included in the index.
<code>type</code>	The type of data source. Specify <code>GITHUB</code> as your data source type.
<code>enableIdentityCrawler</code>	Specify <code>true</code> to use the Amazon Q identity crawler to sync identity/principal information on users and groups with access to specific documents.

Configuration	Description
syncMode	<p>Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose between the following options:</p> <ul style="list-style-type: none"> • Use <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index. • Use <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index. • Use <code>CHANGE_LOG</code> to incrementally crawl only new and modified content each time your data source syncs with your index.
secretArn	<p>The Amazon Resource Name (ARN) of an AWS Secrets Manager secret that contains the key-value pairs required to connect to your GitHub (Server). The secret must contain a JSON structure with the following keys:</p> <pre data-bbox="829 1291 1507 1451"> { "personalToken": " <i>token</i>" } </pre>
version	<p>The version of this template that's currently supported.</p>

How Amazon Q Business connector crawls GitHub (Server) ACLs

When you connect an GitHub (Server) data source to Amazon Q Business, Amazon Q crawls ACL information attached to a document (user and group information) from your GitHub (Server)

instance. If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

The user IDs are mapped as follows:

- `_user_id` – User IDs exist in GitHub on files where there are set access permissions. They are mapped from the user emails as the IDs in GitHub.

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business GitHub (Server) data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q Business enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a attribute mapping already available, or if you want to map additional document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

⚠ Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q GitHub connector supports the following entities and the associated reserved and custom attributes.

📘 Note

You can map any GitHub field to the document title or document body Amazon Q reserved/default index fields.

Supported entities and field mappings

- [Repository](#)
- [Repository Commit](#)
- [Issue Document](#)
- [Issue Comment](#)
- [Issue Attachment](#)
- [Pull Request Comment](#)
- [Pull Request Document](#)
- [Pull Request Attachment](#)

Repository

GitHub field name	Index field name	Description	Data type
Description	_document_body	Default	String
repositoryName	gh_repository_name	Custom	String
repositoryVisibility	gh_repository_visibility	Custom	String

GitHub field name	Index field name	Description	Data type
category	_category	Default	String
owner	_authors	Default	String list
sourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date

Repository Commit

GitHub field name	Index field name	Description	Data type
Description	_document_body	Default	String
repositoryName	gh_repository_name	Custom	String
repositoryVisibility	gh_repository_visibility	Custom	String
category	_category	Default	String
fileType	_file_type	Default	String
owner	_authors	Default	String list
sourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date
fileName	gh_file_name	Default	String
fileSize	gh_size	Default	Long (numeric)
branchName	gh_branch_name	Default	String

Issue Document

GitHub field name	Index field name	Description	Data type
repositoryName	gh_repository_name	Custom	String
repositoryVisibility	gh_repository_visibility	Custom	String
category	_category	Default	String
issueNumber	gh_issue_number	Custom	Long (numeric)
issueTitle	gh_issue_title	Custom	String
owner	_authors	Default	String list
fileType	_file_type	Default	String
issueSourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date
issueFileName	gh_file_name	Custom	String
issueState	gh_issue_state	Custom	String
issueLabel	gh_issue_labels	Default	String list
issueAssignee	gh_issue_assignee	Default	String list

Issue Comment

GitHub field name	Index field name	Description	Data type
repositoryName	gh_repository_name	Custom	String

GitHub field name	Index field name	Description	Data type
repositoryVisibility	gh_repository_visibility	Custom	String
category	_category	Default	String
issueNumber	gh_issue_number	Custom	Long (numeric)
issueTitle	gh_issue_title	Custom	String
owner	_authors	Default	String list
issueSourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date
issueState	gh_issue_state	Custom	String
issueLabel	gh_issue_labels	Default	String list
issueAssignee	gh_issue_assignee	Default	String list

Issue Attachment

GitHub field name	Index field name	Description	Data type
repositoryName	gh_repository_name	Custom	String
repositoryVisibility	gh_repository_visibility	Custom	String
category	_category	Default	String
issueNumber	gh_issue_number	Custom	Long (numeric)
issueTitle	gh_issue_title	Custom	String

GitHub field name	Index field name	Description	Data type
owner	_authors	Default	String list
issueSourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date
issueFileName	gh_file_name	Custom	String
issueFileType	_file_type	Custom	String
issueState	gh_issue_state	Custom	String
issueLabel	gh_issue_labels	Default	String list
issueAssignee	gh_issue_assignee	Default	String list

Pull Request Comment

GitHub field name	Index field name	Description	Data type
repositoryName	gh_repository_name	Custom	String
repositoryVisibility	gh_repository_visibility	Custom	String
category	_category	Default	String
PRNumber	gh_pr_number	Custom	Long (numeric)
PRTitle	gh_pr_title	Custom	String
owner	_authors	Default	String list
PRSourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date

GitHub field name	Index field name	Description	Data type
updatedAt	_last_updated_at	Default	Date
PRState	gh_pr_state	Custom	String
PRLabel	gh_pr_labels	Default	String list
PRAssignee	gh_pr_assignee	Default	String list

Pull Request Document

GitHub field name	Index field name	Description	Data type
repositoryName	gh_repository_name	Custom	String
repositoryVisibility	gh_repository_visibility	Custom	String
category	_category	Default	String
PRNumber	gh_number	Custom	Long (numeric)
PRTitle	gh_pr_title	Custom	String
owner	_authors	Default	String list
PRSourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date
PRFileName	gh_file_name	Custom	String
PRFileType	_file_type	Custom	String
PRState	gh_pr_state	Custom	String
PRLabel	gh_pr_labels	Default	String list

GitHub field name	Index field name	Description	Data type
PRAssignee	gh_pr_assignee	Default	String list

Pull Request Attachment

GitHub field name	Index field name	Description	Data type
repositoryName	gh_repository_name	Custom	String
repositoryVisibility	gh_repository_visibility	Custom	String
category	_category	Default	String
PRNumber	gh_number	Custom	Long (numeric)
PRTitle	gh_pr_title	Custom	String
owner	_authors	Default	String list
PRSourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date
PRFileName	gh_file_name	Custom	String
PRFileType	_file_type	Custom	String
PRState	gh_pr_state	Custom	String
PRLabel	gh_pr_labels	Default	String list
PRAssignee	gh_pr_assignee	Default	String list

IAM role for Amazon Q Business GitHub (Server) connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the `BatchPutDocument` and `BatchDeleteDocument` operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- Permission to access the SSL certificate stored in your Amazon S3 bucket.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AllowsAmazonQToGetS3Objects",
    "Action": [
      "s3:GetObject"
    ],
    "Resource": [
      "arn:aws:s3:::{{input_bucket_name}}/*"
    ],
    "Effect": "Allow",
    "Condition": {
      "StringEquals": {
        "aws:ResourceAccount": "{{account_id}}"
      }
    }
  }
],
  {
    "Sid": "AllowsAmazonQToGetSecret",
```

```

    "Effect": "Allow",
    "Action": [
      "secretsmanager:GetSecretValue"
    ],
    "Resource": [
      "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[{{secret_id}}]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToDecryptSecret",
    "Effect": "Allow",
    "Action": [
      "kms:Decrypt"
    ],
    "Resource": [
      "arn:aws:kms:{{region}}:{{account_id}}:key/[{{key_id}}]"
    ],
    "Condition": {
      "StringLike": {
        "kms:ViaService": [
          "secretsmanager.*.amazonaws.com"
        ]
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToIngestDocuments",
    "Effect": "Allow",
    "Action": [
      "qbusiness:BatchPutDocument",
      "qbusiness:BatchDeleteDocument"
    ],
    "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
    {{application_id}}/index/{{index_id}}"
  },
  {
    "Sid": "AllowsAmazonQToIngestPrincipalMapping",
    "Effect": "Allow",
    "Action": [
      "qbusiness:PutGroup",
      "qbusiness:CreateUser",
      "qbusiness>DeleteGroup",
      "qbusiness:UpdateUser",
      "qbusiness:ListGroups"
    ]
  }
}

```

```

    ],
    "Resource": [
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}",
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}",
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}/data-source/*"
    ]
},
{
    "Sid": "AllowsAmazonQToCreateAndDeleteNI",
    "Effect": "Allow",
    "Action": [
        "ec2:CreateNetworkInterface",
        "ec2:DeleteNetworkInterface"
    ],
    "Resource": [
        "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
        "arn:aws:ec2:{{region}}:{{account_id}}:security-group/
[{{security_group}}]"
    ]
},
{
    "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
    "Effect": "Allow",
    "Action": [
        "ec2:CreateNetworkInterface",
        "ec2:DeleteNetworkInterface"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
        "StringLike": {
            "aws:RequestTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
        },
        "ForAllValues:StringEquals": {
            "aws:TagKeys": [
                "AMAZON_Q"
            ]
        }
    }
},
{

```

```

        "Sid": "AllowsAmazonQToCreateTags",
        "Effect": "Allow",
        "Action": [
            "ec2:CreateTags"
        ],
        "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
        "Condition": {
            "StringEquals": {
                "ec2:CreateAction": "CreateNetworkInterface"
            }
        }
    },
    {
        "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
        "Effect": "Allow",
        "Action": [
            "ec2:CreateNetworkInterfacePermission"
        ],
        "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
        "Condition": {
            "StringLike": {
                "aws:ResourceTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
            }
        }
    },
    {
        "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
        "Effect": "Allow",
        "Action": [
            "ec2:DescribeNetworkInterfaces",
            "ec2:DescribeAvailabilityZones",
            "ec2:DescribeNetworkInterfaceAttribute",
            "ec2:DescribeVpcs",
            "ec2:DescribeRegions",
            "ec2:DescribeNetworkInterfacePermissions",
            "ec2:DescribeSubnets"
        ],
        "Resource": "*"
    }
]
}

```


To allow Amazon Q to assume a role, you must also use the following trust policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToAssumeRoleForServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "{{source_account}}"
        },
        "ArnLike": {
          "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
        }
      }
    }
  ]
}
```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Connecting Gmail to Amazon Q Business

Gmail is an email client developed by Google through which you can send email messages with file attachments. Gmail messages can be sorted and stored inside your email inbox using folders and labels. You can connect Gmail instance to Amazon Q Business—using either the AWS Management Console or the [CreateDataSource](#) API—and create an Amazon Q web experience.

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).

- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [Gmail connector overview](#)
- [Prerequisites for connecting Amazon Q Business to Gmail](#)
- [Connecting Amazon Q Business to Gmail using the console](#)
- [Connecting Amazon Q Business to Gmail using APIs](#)
- [How Amazon Q Business connector crawls Gmail ACLs](#)
- [Amazon Q Business Gmail data source connector field mappings](#)
- [IAM role for Amazon Q Business Gmail connector](#)
- [Troubleshooting your Amazon Q Business Gmail connector](#)

Gmail connector overview

The following table gives an overview of the Amazon Q Business Gmail connector and its supported features.

Category	Feature	Support
Security	Authentication type	Google Service Account
	Authentication credentials	Google service account <ul style="list-style-type: none"> • Google service account • Admin account email • Client email • Private key
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Identity crawling	No
	VPC	Yes

Category	Feature	Support
Crawl features	Custom metadata	No
	Entities	Yes. The following entities are supported: <ul style="list-style-type: none"> • Messages • Attachments
	Field mappings	Yes. Supports default field mappings. For more information, see Field mappings .
	Filters	Yes. The following filters are supported: <ul style="list-style-type: none"> • Include/ exclude message attachments • Include content by date range. • Include/exclude content by email from, to, cc, and bcc domains • Include/exclude content by keywords in subjects • Include/exclude content by label name • Include/exclude content by file name and file type
	Sync mode	Supports full and incremental sync.
	File types	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to Gmail

Before you begin, make sure that you have completed the following prerequisites.

In Gmail, make sure you have:

- Created a Google Cloud Platform admin account and have created a Google Cloud project.
- Activated the Gmail API and Admin SDK API in your admin account.

- Created a service account and downloaded a JSON private key for your Gmail. For information about how to create and access your private key, see [Create a service account key](#) and [Service account credentials](#) on the Google Cloud website.
- Copied your admin account email, your service account email, and your private key to use for authentication.
- Added the following OAuth scopes (using an admin role) for your user and the shared directories you want to index:
 - <https://www.googleapis.com/auth/admin.directory.user.readonly>
 - <https://www.googleapis.com/auth/gmail.readonly>

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your Gmail authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Connecting Amazon Q Business to Gmail using the console

The following procedure outlines how to connect Amazon Q Business to Gmail using the AWS Management Console.

Connecting Amazon Q to Gmail

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).

4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **Gmail** page, enter the following information:
6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

 **Note**

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

8. In **Authentication**, for **AWS Secrets Manager secret** – Choose an existing secret or create a Secrets Manager secret to store your Gmail authentication credentials. If you choose to create a secret, an AWS Secrets Manager secret window opens.
 - Enter the following information in the **Create an AWS Secrets Manager secret window**:
 - i. **Secret Name** – A name for your secret.
 - ii. **Client email** – The client email address that you copied from your Google service account.
 - iii. **Admin account email** – The admin account email address that you would like to use.
 - iv. **Private key** – The private key that you copied from your Google service account.
 - v. Choose **Save**.
9. **Configure VPC and security group** – *optional* – Choose whether you want to use a VPC. If you do, enter the following information:
 - a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.

- b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

10. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

11. In **Sync scope**, for **Entity types** – Choose if you want to crawl **Message attachments**. Messages are crawled by default.
12. For **Additional configuration – optional**, enter the following information:

- a. **Date range** – Enter a date range to specify the start and end date of email messages to be crawled.
- b. **Email domains** – Include or exclude email messages based on domains.
- c. **Keywords in subjects** – Include or exclude email messages based on keywords in their subjects.

 **Note**

You can also choose to include any documents that match all the subject keywords that you have entered.

- d. **Labels** – Add regular expression patterns to include or exclude specific labels. You can add up to 100 patterns.
 - e. **Attachments** – Add regular expression patterns to include or exclude specific attachments. You can add up to 100 patterns.
13. For **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.
 - **Full sync** – Sync all content regardless of the previous sync status.
 - **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.
 14. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).

15. **Tags - *optional*** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
16. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:
 - a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
 - b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

 **Note**

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

17. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

18. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

 **Note**

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to Gmail using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the configuration parameter to provide a JSON schema with all other configuration information specific to your data source connector.

For an example of the API request, see [CreateDataSource](#) in the Amazon Q API Reference.

Gmail JSON schema

The following is the Gmail JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
      "type": "object",
      "properties": {
      }
    },
    "repositoryConfigurations": {
      "type": "object",
      "properties": {
        "message": {
          "type": "object",
          "properties": {
            "fieldMappings": {
              "type": "array",
              "items": [
                {
                  "type": "object",
                  "properties": {
                    "indexFieldName": {
                      "type": "string"
                    },
                    "indexFieldType": {
                      "type": "string",
                      "enum": ["STRING", "STRING_LIST", "DATE"]
                    },
                    "dataSourceFieldName": {
                      "type": "string"
                    },
                    "dateFieldFormat": {
                      "type": "string"
                    }
                  }
                }
              ]
            }
          }
        }
      }
    }
  }
}
```



```

        "required": [
            "indexFieldName",
            "indexFieldType",
            "dataSourceFieldName"
        ]
    }
]
}
},
"attachments": {
    "type": "object",
    "properties": {
        "fieldMappings": {
            "type": "array",
            "items": [
                {
                    "type": "object",
                    "properties": {
                        "indexFieldName": {
                            "type": "string"
                        },
                        "indexFieldType": {
                            "type": "string",
                            "enum": ["STRING"]
                        },
                        "dataSourceFieldName": {
                            "type": "string"
                        }
                    },
                    "required": [
                        "indexFieldName",
                        "indexFieldType",
                        "dataSourceFieldName"
                    ]
                }
            ]
        }
    }
},
"required": []
},
"additionalProperties": {

```

```
"type": "object",
"properties": {
  "isCrawlAcl": {
    "type": "boolean"
  },
  "fieldForUserId": {
    "type": "string"
  },
  "inclusionLabelNamePatterns": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "exclusionLabelNamePatterns": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "inclusionAttachmentTypePatterns": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "exclusionAttachmentTypePatterns": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "inclusionAttachmentNamePatterns": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "exclusionAttachmentNamePatterns": {
    "type": "array",
    "items": {
      "type": "string"
    }
  }
},
```

```
"inclusionSubjectFilter": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"exclusionSubjectFilter": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"isSubjectAnd": {
  "type": "boolean"
},
"inclusionFromFilter": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"exclusionFromFilter": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"inclusionToFilter": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"exclusionToFilter": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"inclusionCcFilter": {
  "type": "array",
  "items": {
    "type": "string"
  }
}
```

```
    },
    "exclusionCcFilter": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    "inclusionBccFilter": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    "exclusionBccFilter": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    "beforeDateFilter": {
      "anyOf": [
        {
          "type": "string",
          "pattern": "^[0-9]{4}-[0-9]{2}-[0-9]{2}T[0-9]{2}:[0-9]{2}:[0-9]{2}Z$"
        },
        {
          "type": "string",
          "pattern": ""
        }
      ]
    },
    "afterDateFilter": {
      "anyOf": [
        {
          "type": "string",
          "pattern": "^[0-9]{4}-[0-9]{2}-[0-9]{2}T[0-9]{2}:[0-9]{2}:[0-9]{2}Z$"
        },
        {
          "type": "string",
          "pattern": ""
        }
      ]
    },
    "isCrawlAttachment": {
```

```

        "type": "boolean"
    },
    "shouldCrawlDraftMessages": {
        "type": "boolean"
    }
},
"required": [
    "isCrawlAttachment",
    "shouldCrawlDraftMessages"
]
},
"type" : {
    "type" : "string",
    "pattern": "GMAIL"
},
"syncMode": {
    "type": "string",
    "enum": [
        "FORCED_FULL_CRAWL",
        "FULL_CRAWL"
    ]
},
"enableIdentityCrawler": {
    "type": "boolean"
},
"secretArn": {
    "type": "string"
},
"version": {
    "type": "string",
    "anyOf": [
        {
            "pattern": "1.0.0"
        }
    ]
}
},
"required": [
    "connectionConfiguration",
    "repositoryConfigurations",
    "additionalProperties",
    "syncMode",
    "secretArn",
    "type"

```

```
]
}
```

The following table provides information about important JSON keys to configure.

Configuration	Description
<code>connectionConfiguration</code>	Configuration information for the endpoint for the data source.
<code>repositoryConfigurations</code>	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings. Specify the type of data source and the secret ARN.
<ul style="list-style-type: none"> <code>message</code> <code>attachments</code> 	A list of objects that map the attributes or field names of your Gmail messages and attachments to Amazon Q index field names.
<code>additionalProperties</code>	Additional configuration options for your content in your data source.
<code>isCrawlAcl</code>	Specify <code>true</code> to crawl access control information from documents.
<code>fieldForUserId</code>	Specify field to use for <code>UserId</code> for ACL crawling.
<ul style="list-style-type: none"> <code>inclusionLabelNamePatterns</code> <code>exclusionLabelNamePatterns</code> <code>inclusionAttachmentTypePatterns</code> <code>exclusionAttachmentTypePatterns</code> <code>inclusionAttachmentNamePatterns</code> 	A list of regular expression patterns to include or exclude messages with specific subject names in your Gmail data source. Files that match the patterns are included in the index. If a file matches both an inclusion and an exclusion pattern, the exclusion pattern takes precedence, and the file isn't included in the index.

Configuration	Description
<ul style="list-style-type: none"> • <code>exclusionAttachmentNamePatterns</code> • <code>inclusionSubjectFilter</code> • <code>exclusionSubjectFilter</code> • <code>inclusionFromFilter</code> • <code>exclusionFromFilter</code> • <code>inclusionToFilter</code> • <code>exclusionToFilter</code> • <code>inclusionCcFilter</code> • <code>exclusionCcFilter</code> • <code>inclusionBccFilter</code> • <code>exclusionBccFilter</code> 	
<code>isSubjectAnd</code>	true to index.
<code>beforeDateFilter</code>	Specify messages and attachments to be included before a certain date.
<code>afterDateFilter</code>	Specify messages and attachments to be included after a certain date.
<code>isCrawlAttachment</code>	A Boolean value to choose whether you want to crawl attachments. Messages are automatically crawled.
<code>type</code>	The type of data source. Specify GMAIL as your data source type.
<code>shouldCrawlDraftMessages</code>	A Boolean value to choose whether you want to crawl draft messages.

Configuration	Description
syncMode	<p>Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose from the following options:</p> <ul style="list-style-type: none">• Use <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index• Use <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index <div data-bbox="829 898 1507 1749" style="border: 1px solid #f08080; border-radius: 10px; padding: 10px;"><p>⚠ Important</p><p>Because there is no API to update permanently deleted Gmail messages, a New, modified, or deleted content sync does <i>not</i> do the following:</p><ul style="list-style-type: none">• Remove messages that were permanently deleted from Gmail from your Amazon Q index• Sync changes in Gmail email labels<p>To sync your Gmail data source label changes and permanently deleted email messages to your Amazon Q index, you must run full crawls periodically.</p></div>

Configuration	Description
enableIdentityCrawler	Specify true to use the Amazon Q identity crawler to sync identity/principal information on users and groups with access to specific documents.
secretARN	<p>The Amazon Resource Name (ARN) of a Secrets Manager secret that contains the key-value pairs required to connect to your Gmail. The secret must contain a JSON structure with the following keys:</p> <pre data-bbox="829 716 1507 1031"> { "adminAccountEmailId": "\${adminAccountEmailId}" , "clientEmailId": "\${clientEmailId}" , "privateKey": "\${privateKey}" } </pre>
version	The version of the template that's currently supported.

How Amazon Q Business connector crawls Gmail ACLs

When you connect an Gmail data source to Amazon Q Business, Amazon Q Business crawls ACL information attached to a document (user and group information) from your Gmail instance. If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

The user IDs are mapped as follows:

- `_user_id` – User IDs exist in Gmail on files where there are set access permissions. They're mapped from the user emails as the IDs in Gmail.

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business Gmail data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q Business enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a attribute mapping already available, or if you want to map additional document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q Gmail connector supports the following entities and the associated reserved and custom attributes.

Supported entities and field mappings

- [Messages](#)

Messages

Gmail field name	Index field name	Description	Data type
category	_category	Default	String
internalDate	_created_at	Default	Date
id	gmail_message_is	Custom	String
labelIds	gmail_message_label_ids	Custom	String list
historyId	gmail_message_history_id	Custom	String
subject	gmail_subject	Custom	String
from	gmail_from	Custom	String
to	gmail_to	Custom	String list
cc	gmail_cc	Custom	String list
bcc	gmail_bcc	Custom	String list

IAM role for Amazon Q Business Gmail connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the `BatchPutDocument` and `BatchDeleteDocument` operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToGetSecret",
      "Effect": "Allow",
      "Action": [
        "secretsmanager:GetSecretValue"
      ],
      "Resource": [
        "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
      ]
    },
    {
      "Sid": "AllowsAmazonQToDecryptSecret",
      "Effect": "Allow",
      "Action": [
        "kms:Decrypt"
      ],
      "Resource": [
        "arn:aws:kms:{{region}}:{{account_id}}:key/[key_id]"
      ],
      "Condition": {
        "StringLike": {
          "kms:ViaService": [
            "secretsmanager.*.amazonaws.com"
          ]
        }
      }
    },
    {
      "Sid": "AllowsAmazonQToIngestDocuments",
      "Effect": "Allow",
```

```

    "Action": [
      "qbusiness:BatchPutDocument",
      "qbusiness:BatchDeleteDocument"
    ],
    "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
  },
  {
    "Sid": "AllowsAmazonQToIngestPrincipalMapping",
    "Effect": "Allow",
    "Action": [
      "qbusiness:PutGroup",
      "qbusiness:CreateUser",
      "qbusiness:DeleteGroup",
      "qbusiness:UpdateUser",
      "qbusiness:ListGroups"
    ],
    "Resource": [
      "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}",
      "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}",
      "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}/data-source/*"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateAndDeleteNI",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2:DeleteNetworkInterface"
    ],
    "Resource": [
      "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
      "arn:aws:ec2:{{region}}:{{account_id}}:security-group/[{{security_group}}]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2:DeleteNetworkInterface"
    ],
  },

```

```

"Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
"Condition": {
  "StringLike": {
    "aws:RequestTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
  },
  "ForAllValues:StringEquals": {
    "aws:TagKeys": [
      "AMAZON_Q"
    ]
  }
}
},
{
  "Sid": "AllowsAmazonQToCreateTags",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateTags"
  ],
  "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
  "Condition": {
    "StringEquals": {
      "ec2:CreateAction": "CreateNetworkInterface"
    }
  }
},
{
  "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateNetworkInterfacePermission"
  ],
  "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
  "Condition": {
    "StringLike": {
      "aws:ResourceTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
    }
  }
},
{
  "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
  "Effect": "Allow",
  "Action": [
    "ec2:DescribeNetworkInterfaces",
    "ec2:DescribeAvailabilityZones",

```

```

    "ec2:DescribeNetworkInterfaceAttribute",
    "ec2:DescribeVpcs",
    "ec2:DescribeRegions",
    "ec2:DescribeNetworkInterfacePermissions",
    "ec2:DescribeSubnets"
  ],
  "Resource": "*"
}
]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "{{source_account}}"
        },
        "ArnEquals": {
          "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
        }
      }
    }
  ]
}

```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Troubleshooting your Amazon Q Business Gmail connector

The following table provides information about error codes you may see for the Gmail connector and suggested troubleshooting actions.

Error code	Error message	Suggested resolution
GML-5001	There was a problem while retrieving directory .	There was a problem while retrieving directory because of incorrect credentials. Provide correct credentials and try again.
GML-5002	There was a problem while retrieving user specific Gmail object.	There was a problem while retrieving user specific Gmail object because of incorrect credentials. Provide correct credentials and try again.
GML-5003	Connection lost - A problem occurred while validating credentials.	Connection was lost due to invalid credentials. Provide correct credentials and try again.
GML-5004	There was a problem while retrieving the user list because the API was not responding.	There was a problem while retrieving the user list because the API was not responding. Try again.
GML-5100	There was a problem while retrieving repository configurations. Repository configurations may be empty or incorrect.	Repository configurations should not be empty or incorrect. Provide valid details for repository configurations.
GML-5101	There was a problem while retrieving message entity from repository configurations. No message entity found in repository configurations.	Message entity should not be empty. Check if message entity is present in repository configurations and provide the same if not present.
GML-5102	There was a problem while retrieving	Attachment entity should not be empty. Check if attachment entity is

Error code	Error message	Suggested resolution
	attachment entity from repository configurations. No attachment entity found in repository configurations.	present in repository configurations and provide the same if not present.
GML-5103	There was a problem while retrieving field mappings for message entity from repository configurations. Field mappings may be empty or incorrect.	Field mappings should not be empty or incorrect. Provide proper field mappings for message entity in repository configurations.
GML-5104	There was a problem while retrieving field mappings for attachment entity from repository configurations. Field mappings may be empty or incorrect.	Field mappings should not be empty or incorrect. Provide proper field mappings for message entity in repository configurations.
GML-5105	There was a problem while retrieving field mapping values for message entity. Field mapping values may be empty or incorrect.	Field mappings values should not be empty or incorrect. Provide proper field mapping values for message entity in repository configurations.
GML-5106	There was a problem while retrieving field mapping values for attachment entity. Field mapping values may be empty or incorrect.	Field mappings values should not be empty or incorrect. Provide proper field mapping values for message entity in repository configurations.

Error code	Error message	Suggested resolution
GML-5107	There was a problem while parsing before/after date filter value. Before/After date format may be incorrect.	Provide correct before/after date format. E.g. yyyy-MM-ddTHH:mm:ssZ.
GML-5108	There was a problem while retrieving client email id. Client email id may be empty or incorrect.	The client email id should not be empty or incorrect. Provide correct client email id.
GML-5109	There was a problem while retrieving admin account email id. Admin account email id may be empty or incorrect.	The admin account email id should not be empty or incorrect. Provide correct admin account email id.
GML-5110	There was a problem while retrieving private key. Private key may be empty or incorrect.	The private key should not be empty or incorrect. Provide correct private key.
GML-5111	One or more of the provided filter regex are invalid.	Provide correct regex value in filter fields.
GML-5200	There was a problem while retrieving Gmail items.	There was a problem while retrieving Gmail items because user is not provided. Ensure that user is not empty.

Error code	Error message	Suggested resolution
GML-5201	There was a problem while retrieving the message body because the API was not responding.	There was a problem while retrieving the message body because the API was not responding. Try again.
GML-5202	There was a problem while retrieving the message subject because the API was not responding.	There was a problem while retrieving the message subject because the API was not responding. Try again.
GML-5203	There was a problem while retrieving the attachment because the API was not responding.	There was a problem while retrieving the attachment because the API was not responding. Try again.
GML-5204	There was a problem while retrieving the message metadata because the API was not responding.	There was a problem while retrieving the message metadata because the API was not responding. Try again.
GML-5205	There was a problem while retrieving the attachment metadata because the API was not responding.	There was a problem while retrieving the attachment metadata because the API was not responding. Try again.
GML-5206	There was a problem while retrieving the message because the API was not responding.	There was a problem while retrieving the message because the API was not responding. Try again.

Error code	Error message	Suggested resolution
GML-5500	Connection timed out - API is not responding. The threshold number of API calls has been exceeded.	Timeout exception occurred due to API not responding. The threshold number of API hits has been exceeded. Try again.

Connecting Google Drive to Amazon Q Business

Google Drive is a cloud-based file storage service. Amazon Q Business can connect to your Google Drive instances. You can connect Google Drive instance to Amazon Q—using either the AWS Management Console or the [CreateDataSource](#) API—and create an Amazon Q web experience.

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).
- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [Google Drive connector overview](#)
- [Prerequisites for connecting Amazon Q Business to Google Drive](#)
- [Connecting Amazon Q Business to Google Drive using the console](#)
- [Connecting Amazon Q Business to GoogleDrive using APIs](#)
- [How Amazon Q Business connector crawls GoogleDrive ACLs](#)
- [Amazon Q Business Google Drive data source connector field mappings](#)
- [IAM role for Amazon Q Business Google Drive connector](#)
- [Known limitations for the Amazon Q Business Google Drive connector](#)
- [Troubleshooting your Amazon Q Business Google Drive connector](#)

Google Drive connector overview

The following table gives an overview of the Amazon Q Business Google Drive connector and its supported features.

Category	Feature	Support
Security	Authentication type	Google Service Account, OAuth 2.0
	Authentication credentials	<p>Google service account</p> <ul style="list-style-type: none"> Admin account email Client email Private key <p>OAuth 2.0</p> <ul style="list-style-type: none"> Client ID Client secret Refresh token <div style="border: 1px solid #f08080; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p>⚠ Important Admin privileges required.</p> </div>
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Identity crawling	Yes. Supported only with Google service account authentication.
	VPC	Yes
Crawl features	Custom metadata	No
	Entities	Yes. The following entities are supported: <ul style="list-style-type: none"> Files

Category	Feature	Support
		<ul style="list-style-type: none"> • Comments
	Field mappings	Yes. Supports default field mappings. For more information, see Field mappings .
	Filters	<p>Yes. The following filters are supported:</p> <ul style="list-style-type: none"> • Include files based on file size • Include/exclude Shared drives • Include/exclude by mime types • Include/exclude attachments by file name, file type, and file path
	Sync mode	Supports full and incremental sync.
	File types	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to Google Drive

Before you begin, make sure that you have completed the following prerequisites.

In Google Drive, make sure you have:

- **Either** been granted access by a super admin role **or** are a user with administrative privileges. You do not need a super admin role for yourself if you have been granted access by a super admin role.
- Configured Google Drive Service Account connection credentials containing your admin account email, client email (service account email), and private key. See [Google Cloud documentation on creating and deleting service account keys](#).
- Created a Google Cloud Service Account (an account with delegated authority to assume a user identity) with **Enable G Suite Domain-wide Delegation** activated for server-to-server authentication, and then generated a JSON private key using the account.

Note

The private key should be generated after the creation of the service account.

- Added Admin SDK API and Google Drive API in your user account.
- **Optional:** Configured Google Drive OAuth 2.0 connection credentials containing client ID, client secret, and refresh token as connection credentials for a specific user. You need this to crawl individual account data. See [Google documentation on using OAuth 2.0 to access APIs](#).
- Added (or asked a user with a super admin role to add) the following OAuth scopes to your service account using a super admin role. These API scopes are needed to crawl all documents, and access control (ACL) information for all users in a Google Workspace domain:
 - <https://www.googleapis.com/auth/drive.readonly>—View and download all your Google Drive files
 - <https://www.googleapis.com/auth/drive.metadata.readonly>—View metadata for files in your Google Drive
 - <https://www.googleapis.com/auth/admin.directory.group.readonly>—Scope for only retrieving group, group alias, and member information. This is needed for the Amazon Q Identity Crawler.
 - <https://www.googleapis.com/auth/admin.directory.user.readonly>—Scope for only retrieving users or user aliases. This is needed for listing users in the Amazon Q Identity Crawler and for setting ACLs.
 - <https://www.googleapis.com/auth/cloud-platform>—Scope for generating access token for fetching content of large Google Drive files.
 - <https://www.googleapis.com/auth/forms.body.readonly>—Scope for fetching data from Google Forms.

To support the Forms API, add the following additional scope:

- <https://www.googleapis.com/auth/forms.body.readonly>

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.

- Stored your Google Drive authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

 **Note**

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Connecting Amazon Q Business to Google Drive using the console

The following procedure outlines how to connect Amazon Q Business to Google Drive using the AWS Management Console.

Connecting Amazon Q to Google Drive

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **Google Drive** page, enter the following information:
6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

Note

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

8. For **Authentication** – Choose between **Google service account** and **OAuth 2.0 authentication**, based on your use case.
9. **AWS Secrets Manager secret** – Choose an existing secret or create a Secrets Manager secret to store your GoogleDrive authentication credentials. If you choose to create a secret, an AWS Secrets Manager secret window opens.

- If you choose **Existing**, select an existing secret for **Select secret**.

If you choose **New**, enter the following information in the **New AWS Secrets Manager secret** section:

- i. **Secret name** – A name for your secret.
- ii. If you chose **Google service account**, enter the following information:
 - **Secret Name** – A name for your secret.
 - **Admin account email** – The email ID of the admin user (the email used by the Service Account User) in your Google service account configuration.
 - **Client email** – The email ID of the service account.
 - **Private Key** – The private key created in your service account.

Then, choose **Save and add secret**.

- iii. If you chose **OAuth 2.0 authentication**, enter the details of **Secret Name**, **Client ID**, **Client secret** and **Refresh token** that you created in your service account. Then, choose **Save and add secret**.

10. **Configure VPC and security group** – *optional* – Choose whether you want to use a VPC. If you do, enter the following information:
 - a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.

- b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

11. **Identity crawler** – Choose to activate Amazon Q identity crawler to sync identity information. For more information, see [Identity crawler](#).
12. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

13. In **Sync scope**, for **Sync contents** – Choose from the following options to select content to index:

 **Note**

To further limit content to index, use **Entity regex patterns** in the **Additional configuration** section.

- **My Drive & Shared with me – My Drive** contains a user's personal folders and documents. **Shared with me** contains all the folders and documents that have been shared with the user. Select this option to index both.
 - **Shared drives – Shared drives** are folders used to store, access, and share files with a team. Select this option to index these.
 - **Comments** – Select this option to index file comments.
14. In **Additional configuration - optional**, enter the following optional information:
 - a. **Maximum file size** – Set the maximum file size value that Amazon Q will crawl.
 - b. **User email** – Add the user email IDs that you want to include or exclude.
 - c. **Shared drives** – Add the shared drives that you want to include or exclude.
 - d. **Mime types** – Add the MIME types that you want to include or exclude.
 - e. **Entity patterns** – Add regular expression patterns to include or exclude certain folders, files, and file types from **My drive**, **Shared with me**, and **Shared drives**. You can add up to 100 patterns.

15. In **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.
- **Full sync** – Sync all content regardless of the previous sync status.
 - **New or modified content sync** – Sync only new and modified documents.
 - **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.

For more details, see [Sync mode](#).

16. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
17. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
18. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:
- a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
 - b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

 **Note**

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

19. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

20. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

Note

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to GoogleDrive using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the `configuration` parameter to provide a JSON schema with all other configuration information specific to your data source connector.

For an example of the API request, see [CreateDataSource](#) in the Amazon Q API Reference.

Google Drive JSON schema

The following is the Google Drive JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
      "type": "object",
      "properties": {
        "repositoryEndpointMetadata": {
          "type": "object",
          "properties": {
            "authType": {
              "type": "string",
              "enum": [
                "serviceAccount",
                "OAuth2"
              ]
            }
          }
        },
        "required": [
          "authType"
        ]
      }
    }
  }
}
```

```
    }
  },
  "required": [
    "repositoryEndpointMetadata"
  ]
},
"repositoryConfigurations": {
  "type": "object",
  "properties": {
    "file": {
      "type": "object",
      "properties": {
        "fieldMappings": {
          "type": "array",
          "items": [
            {
              "type": "object",
              "properties": {
                "indexFieldName": {
                  "type": "string"
                },
                "indexFieldType": {
                  "type": "string",
                  "enum": [
                    "STRING",
                    "DATE",
                    "STRING_LIST",
                    "LONG"
                  ]
                },
                "dataSourceFieldName": {
                  "type": "string"
                },
                "dateFieldFormat": {
                  "type": "string",
                  "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
                }
              }
            }
          ]
        },
        "required": [
          "indexFieldName",
          "indexFieldType",
          "dataSourceFieldName"
        ]
      }
    }
  }
}
```

```

    ]
  }
},
"required": [
  "fieldMappings"
]
},
"comment": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": [
        {
          "type": "object",
          "properties": {
            "indexFieldName": {
              "type": "string"
            },
            "indexFieldType": {
              "type": "string",
              "enum": [
                "STRING",
                "DATE",
                "STRING_LIST"
              ]
            },
            "dataSourceFieldName": {
              "type": "string"
            },
            "dateFieldFormat": {
              "type": "string",
              "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
            }
          }
        },
        "required": [
          "indexFieldName",
          "indexFieldType",
          "dataSourceFieldName"
        ]
      ]
    }
  }
},

```

```
        "required": [
          "fieldMappings"
        ]
      }
    },
    "additionalProperties": {
      "type": "object",
      "properties": {
        "maxFileSizeInMegabytes": {
          "type": "string"
        },
        "isCrawlComment": {
          "type": "boolean"
        },
        "isCrawlMyDriveAndSharedWithMe": {
          "type": "boolean"
        },
        "isCrawlSharedDrives": {
          "type": "boolean"
        },
        "isCrawlAcl": {
          "type": "boolean"
        },
        "fieldForUserId": {
          "type": "string"
        },
        "excludeUserAccounts": {
          "type": "array",
          "items": {
            "type": "string"
          }
        },
        "excludeSharedDrives": {
          "type": "array",
          "items": {
            "type": "string"
          }
        },
        "excludeMimeType": {
          "type": "array",
          "items": {
            "type": "string"
          }
        }
      }
    }
  }
}
```

```
  },
  "includeUserAccounts": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "includeSharedDrives": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "includeMimeType": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "includeTargetAudienceGroup": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "inclusionFileTypePatterns": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "inclusionFileNamePatterns": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "exclusionFileTypePatterns": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "exclusionFileNamePatterns": {
```



```
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "inclusionFilePathFilter": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "exclusionFilePathFilter": {
    "type": "array",
    "items": {
      "type": "string"
    }
  }
},
"type": {
  "type": "string",
  "pattern": "GOOGLEDRIVEV2"
},
"enableIdentityCrawler": {
  "type": "boolean"
},
"syncMode": {
  "type": "string",
  "enum": [
    "FORCED_FULL_CRAWL",
    "FULL_CRAWL",
    "CHANGE_LOG"
  ]
},
"secretArn": {
  "type": "string",
  "minLength": 20,
  "maxLength": 2048
}
},
"version": {
  "type": "string",
  "anyOf": [
    {
```

```

    "pattern": "1.0.0"
  }
]
},
"required": [
  "connectionConfiguration",
  "repositoryConfigurations",
  "syncMode",
  "additionalProperties",
  "secretArn",
  "type"
]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the data source.
repositoryEndpointMetadata	The endpoint information for the data source. This data source doesn't specify an endpoint. You choose your authentication type: <code>serviceAccount</code> and <code>OAuth2</code> . The connection information is included in an AWS Secrets Manager secret that you provide the <code>secretArn</code> .
authType	Choose between <code>serviceAccount</code> and <code>OAuth2</code> , based on your use case.
repositoryConfigurations	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings.
<ul style="list-style-type: none"> file comment 	A list of objects that map the attributes or field names of your Google Drive to Amazon Q index field names.

Configuration	Description
<code>additionalProperties</code>	Additional configuration options for your content in your data source
<code>isCrawlAcl</code>	Specify <code>true</code> to crawl access control information from documents.
<code>fieldForUserId</code>	Specify field to use for <code>UserId</code> for ACL crawling.
<ul style="list-style-type: none"> <code>maxFileSizeInMegabytes</code> 	Enter the maximum file size value in MBs that Amazon Q should crawl and index from your Google Drive data source.
<ul style="list-style-type: none"> <code>iscrawlComment</code> 	<code>true</code> to index comments in your Google Drive data source.
<ul style="list-style-type: none"> <code>isCrawlMyDriveAndSharedWithMe</code> 	<code>true</code> to index MyDrive and Shared With Me Drives in your Google Drive data source.
<ul style="list-style-type: none"> <code>isCrawlSharedDrives</code> 	<code>true</code> to index Shared Drives in your Google Drive data source.
<ul style="list-style-type: none"> <code>excludeUserAccounts</code> <code>excludeSharedDrives</code> <code>excludeMimeType</code> <code>exclusionFileTypePatterns</code> <code>exclusionFileNamePatterns</code> <code>exclusionFilePathFilter</code> 	A list of regular expression patterns to <i>exclude</i> specific files in your Google Drive data source. Files that match the patterns are excluded from the index. Files that don't match the patterns are included in the index. If a file matches both an exclusion and inclusion pattern, the exclusion pattern takes precedence, and the file isn't included in the index.

Configuration	Description
<ul style="list-style-type: none">• <code>includeUserAccounts</code>• <code>includeSharedDrives</code>• <code>includeMimeType</code>• <code>inclusionFileTypePatterns</code>• <code>inclusionFileNamePatterns</code>• <code>inclusionFilePathFilter</code>	A list of regular expression patterns to <i>include</i> specific files in your Google Drive data source. Files that match the patterns are included in the index. Files that don't match the patterns are excluded from the index. If a file matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence, and the file isn't included in the index.
<code>type</code>	The type of data source. Specify <code>G000GLEDR IVEV2</code> as your data source type.
<code>enableIdentityCrawler</code>	<code>true</code> to activate identity crawler. Identity crawler is activated by default. Crawling identity information on users and groups with access to certain documents is useful for user context filtering. Search results are filtered based on the user or their group access to documents. See Identity crawler for more information.

Configuration	Description
syncMode	<p>Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose between the following options:</p> <ul style="list-style-type: none">• Use <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index• Use <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index• Use <code>CHANGE_LOG</code> to incrementally crawl only new and modified content each time your data source syncs with your index

Configuration	Description
secretARN	<p>The Amazon Resource Name (ARN) of an AWS Secrets Manager secret that contains the key-value pairs required to connect to your Google Drive. The secret must contain a JSON structure with the following keys:</p> <p>If using Google Service Account authentication:</p> <pre data-bbox="829 617 1507 932"> { "clientEmail": "user account email", "adminAccountEmail": "service account email", "privateKey": "private key" } </pre> <p>If using OAuth 2.0 authentication:</p> <pre data-bbox="829 1045 1507 1276"> { "clientId": "OAuth client ID", "clientSecret": "client secret", "refreshToken": "refresh token" } </pre>
version	The version of this template that's currently supported.

How Amazon Q Business connector crawls GoogleDrive ACLs

When you connect an GoogleDrive data source to Amazon Q Business, Amazon Q Business crawls ACL information attached to a document (user and group information) from your GoogleDrive instance. If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

The GoogleDrive group and user IDs are mapped as follows:

A Google Workspace Drive data source returns user and group information for Google Drive users and groups. Group and domain membership are mapped to the `_group_ids` index field. The Google Drive username is mapped to the `_user_id` field.

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business Google Drive data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q Business enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a attribute mapping already available, or if you want to map additional document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q GoogleDrive connector supports the following entities and the associated reserved and custom attributes.

Supported entities and field mappings

- [Files](#)
- [Comments](#)

Files

GoogleDrive field name	Index field name	Description	Data type
authors	_authors	Default	String list
emailIds	gd_author_emails	Custom	String list
mimeType	gd_file_mime_type	Custom	String
size	gd_size	Custom	Long (numeric)
starred	gd_starred_file	Custom	String
version	gd_size	Custom	Long (numeric)
webViewLink	_source_uri	Default	String
viewedByMeAt	gd_viewed_at	Custom	Date
modifiedByMeAt	gd_modified_by_me_at	Custom	Date
createdAt	_created_at	Default	Date
modifiedAt	_last_updated_at	Default	Date
lastModifyingUser	gd_last_modified_by	Custom	String
kind	gd_kind	Custom	String
id	gd_id	Custom	String

GoogleDrive field name	Index field name	Description	Data type
name	gd_name	Custom	String
parents	gd_parents	Custom	String list
spaces	gd_spaces	Custom	String list
iconLink	gd_icon_link	Custom	String
hasThumbnail	gd_has_thumbnail	Custom	String
thumbnailVersion	gd_thumbnail_version	Custom	Long (numeric)
shared	gd_shared	Custom	String

Comments

GoogleDrive field name	Index field name	Description	Data type
authors	_authors	Default	String list
commentType	gd_type	Custom	String
createdAt	_created_at	Default	Date
modifiedAt	_last_updated_at	Default	Date
webViewLink	_source_uri	Default	String
kind	gd_kind	Custom	String
id	gd_id	Custom	String

IAM role for Amazon Q Business Google Drive connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the `BatchPutDocument` and `BatchDeleteDocument` operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToGetSecret",
      "Effect": "Allow",
      "Action": [
        "secretsmanager:GetSecretValue"
      ],
      "Resource": [
        "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
      ]
    },
    {
      "Sid": "AllowsAmazonQToDecryptSecret",
      "Effect": "Allow",
      "Action": [
        "kms:Decrypt"
      ],
      "Resource": [
```

```

    "arn:aws:kms:{{region}}:{{account_id}}:key/[{{key_id}}]"
  ],
  "Condition": {
    "StringLike": {
      "kms:ViaService": [
        "secretsmanager.*.amazonaws.com"
      ]
    }
  }
},
{
  "Sid": "AllowsAmazonQToIngestDocuments",
  "Effect": "Allow",
  "Action": [
    "qbusiness:BatchPutDocument",
    "qbusiness:BatchDeleteDocument"
  ],
  "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
},
{
  "Sid": "AllowsAmazonQToIngestPrincipalMapping",
  "Effect": "Allow",
  "Action": [
    "qbusiness:PutGroup",
    "qbusiness:CreateUser",
    "qbusiness>DeleteGroup",
    "qbusiness:UpdateUser",
    "qbusiness:ListGroups"
  ],
  "Resource": [
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}",
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}",
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}/data-source/*"
  ]
},
{
  "Sid": "AllowsAmazonQToCreateAndDeleteNI",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateNetworkInterface",
    "ec2>DeleteNetworkInterface"
  ]
}

```

```

    ],
    "Resource": [
      "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
      "arn:aws:ec2:{{region}}:{{account_id}}:security-group/[{{security_group}}]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2:DeleteNetworkInterface"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:RequestTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
      },
      "ForAllValues:StringEquals": {
        "aws:TagKeys": [
          "AMAZON_Q"
        ]
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateTags",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateTags"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringEquals": {
        "ec2:CreateAction": "CreateNetworkInterface"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterfacePermission"
    ],
  },

```

```

"Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
"Condition": {
  "StringLike": {
    "aws:ResourceTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
  }
},
{
  "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
  "Effect": "Allow",
  "Action": [
    "ec2:DescribeNetworkInterfaces",
    "ec2:DescribeAvailabilityZones",
    "ec2:DescribeNetworkInterfaceAttribute",
    "ec2:DescribeVpcs",
    "ec2:DescribeRegions",
    "ec2:DescribeNetworkInterfacePermissions",
    "ec2:DescribeSubnets"
  ],
  "Resource": "*"
}
]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "{{source_account}}"
        },
        "ArnEquals": {
          "aws:SourceArn": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/{{application_id}}"
        }
      }
    }
  ]
}

```

```

    }
  }
}
]
}

```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Known limitations for the Amazon Q Business Google Drive connector

The Amazon Q Google Drive connector has the following known limitations:

- Custom field mapping is not available for Google Drive connector as the Google Drive UI does not support creating custom fields.
- Google Drive API does not support retrieving comments from a permanently deleted file. Comments are retrievable, however, for trashed files. When a file is trashed, the Amazon Q connector will delete comments from the Amazon Q index.
- Google Drive API does not return comments present in a .docx file.

Troubleshooting your Amazon Q Business Google Drive connector

The following table provides information about error codes you may see for the Google Drive connector and suggested troubleshooting actions.

Error code	Error message	Suggested resolution
GD-5001	Invalid directory object message.	There was a problem while retrieving directory.
GD-5002	Invalid user specific object message.	There was a problem while retrieving user specific Drive object.
GD-5003	Error while connecting message.	Could not connect to Google drive - A problem occurred while validating credentials.

Error code	Error message	Suggested resolution
GD-5004	Error fetching user list message.	There was a problem while retrieving the user list because the API was not responding.
GD-5005	Error fetching user file list message.	There was a problem while retrieving the user file list because the API was not responding.
GD-5006	Error fetching user file comment list message.	There was a problem while retrieving the user file comment list because the API was not responding.
GD-5007	Error fetching comment reply list message.	There was a problem while retrieving the comment reply list because the API was not responding.
GD-5008	Error fetching user change list message.	There was a problem while retrieving the user change list because the API was not responding.
GD-5009	Invalid http req initializer message.	There was a problem while retrieving http request initializer.
GD-5010	Invalid new access token message.	There was a problem while generating new access token.
GD-5100	Empty repository configuration message.	There was a problem while retrieving repository configurations. Repository configurations may be empty or incorrect.
GD-5101	Empty file entity message.	There was a problem while retrieving file entity from repository configurations. No file entity found in repository configurations.

Error code	Error message	Suggested resolution
GD-5102	Empty comment entity message.	There was a problem while retrieving comment entity from repository configurations. No comment entity found in repository configurations.
GD-5103	Empty auth type message.	There was a problem while retrieving auth type. Auth type may be empty or incorrect
GD-5104	Empty client id message.	There was a problem while retrieving client id. Client id may be empty or incorrect.
GD-5105	Empty file entity field mapping data message.	There was a problem while retrieving field mapping values for file entity. Field mapping values may be empty or incorrect.
GD-5106	Empty comment entity field mapping data message.	There was a problem while retrieving field mapping values for comment entity. Field mapping values may be empty or incorrect.
GD-5107	Empty client secret message.	There was a problem while retrieving client secret. Client secret may be empty or incorrect.
GD-5114	Empty refresh token message.	There was a problem while retrieving refresh token. Refresh token may be empty or incorrect.
GD-5108	Empty client email message.	There was a problem while retrieving client email id. Client email id may be empty or incorrect.

Error code	Error message	Suggested resolution
GD-5109	Empty client admin account email message.	There was a problem while retrieving client admin account email id. Client admin account email id may be empty or incorrect.
GD-5110	Empty private key message.	There was a problem while retrieving private key. Private key may be empty or incorrect.
GD-5111	Erroneous filter regex.	One or more of the provided filter regex are invalid.
GD-5112	Invalid auth message.	Incorrect auth type. Auth type should be OAuth2 or serviceAccount.
GD-5113	Identity crawler invalid auth message.	Incorrect auth type. Auth type should be serviceAccount.
GD-5115	Invalid user accounts exclusion filter msg.	User accounts for exclusion filter not applicable for OAuth2 auth type.
GD-5116	Invalid user accounts inclusion filter msg.	User accounts for inclusion filter not applicable for OAuth2 auth type.
GD-5200	File content exception.	Exception occurred while fetching File content for file.
GD-5201	Reply content.	Exception occurred while crawling reply.
GD-5202	Comment content.	Exception occurred while crawling comment.
GD-5203	Group fetch.	Exception occurred while crawling group.

Error code	Error message	Suggested resolution
GD-5204	Member fetch.	Exception occurred while crawling member.
GD-5205	File metadata fetch.	Exception occurred while crawling file metadata.
GD-5207	Folder metadata fetch.	Exception occurred while crawling folder metadata.
GD-5208	Drive fetch.	Exception occurred while crawling drive.
GD-5209	Change start token fetch.	Exception occurred while crawling next page token.
GD-5210	Permission list.	Exception occurred while crawling permission list.
GD-5500	Timeout error message.	Connection timed out - API is not responding. The threshold number of API hits has been exceeded.

Connecting IBM DB2 to Amazon Q Business

IBM DB2 is a relational database management system developed by IBM. If you are a AWS user, you can use Amazon Q Business to index your IBM DB2 data source.

You can connect your IBM DB2 instance to Amazon Q—using either the AWS Management Console, CLI, or the [CreateDataSource](#) API—and create an Amazon Q web experience.

The Amazon Q IBM DB2 data source connector supports DB2 11.5.7.

Important

As a best practice, provide Amazon Q with read-only database credentials. Also, avoid adding tables with sensitive data or personal identifiable information (PII).

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).
- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [IBM DB2 connector overview](#)
- [Prerequisites for connecting Amazon Q Business to IBM DB2](#)
- [Connecting Amazon Q Business to IBM DB2 using the console](#)
- [Connecting Amazon Q Business to IBM DB2 using APIs](#)
- [How Amazon Q Business connector crawls IBM DB2 ACLs](#)
- [Amazon Q Business IBM DB2 data source connector field mappings](#)
- [IAM role for Amazon Q Business IBM DB2 connector](#)
- [Known limitations for the Amazon Q Business IBM DB2 connector](#)

IBM DB2 connector overview

The following table gives an overview of the Amazon Q Business IBM DB2 connector and its supported features.

Category	Feature	Support
Security	Authentication type	Basic
	Authentication credentials	<ul style="list-style-type: none"> • Username of database user • Password of database user
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Driver version	DB2 – 11.5.7.0

Category	Feature	Support
	Data source version	11.5.7
	Identity crawling	No
	VPC	Yes
Crawl features	Custom metadata	Yes
	Entities	Yes. The following entities are supported: <ul style="list-style-type: none"> Document <div data-bbox="862 705 1508 974" style="border: 1px solid #add8e6; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p> Note Each database row is considered an individual searchable Amazon Q document.</p> </div>
	Field mappings	Yes. Supports both default and custom field mappings. For more information, see Field mappings .
	Sync mode	Supports full and incremental sync.
	File types	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to IBM DB2

Before you begin, make sure that you have completed the following prerequisites.

In IBM DB2, make sure you have:

- Noted your database user name and password.

Important

As a best practice, provide Amazon Q with read-only database credentials.

- Copied your database host URL, port, and instance.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your IBM DB2 authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Connecting Amazon Q Business to IBM DB2 using the console

The following procedure outlines how to connect Amazon Q Business to IBM DB2 using the AWS Management Console.

Connecting Amazon Q to IBM DB2

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **IBM DB2** page, enter the following information:
6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. In **Source**, enter the following information:

- a. **Host** – Enter the database host name.
 - b. **Port** – Enter the database port.
 - c. **Instance** – Enter the database instance.
 - d. **Enable SSL certificate location** – Choose to enter the Amazon S3 path to your SSL certificate file.
8. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

 **Note**

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

9. In **Authentication** – Enter the following information for your **AWS Secrets Manager secret**.
- a. **Secret name** – A name for your secret.
 - b. For **Database user name**, and **Password** – Enter the authentication credential values you copied from your database.
 - c. Choose **Save**.
10. **Configure VPC and security group** – *optional* – Choose whether you want to use a VPC. If you do, enter the following information:
- a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
 - b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

11. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

12. In **Sync scope**, enter the following information:

- **SQL query** – Enter SQL query statements like SELECT and JOIN operations. SQL queries must be less than 1000 characters and not contain any semi-colons (;). Amazon Q will crawl all database content that matches your query.
- **Primary key column** – Provide the primary key for the database table. This identifies a table within your database.
- **Title column** – Provide the name of the document title column within your database table.
- **Body column** – Provide the name of the document body column within your database table.

13. In **Additional configuration – optional** – Configure the following settings:

- **Change-detecting columns** – Enter the names of the columns that Amazon Q will use to detect content changes. Amazon Q will re-index content when there is a change in any of these columns.
- **Users' IDs column** – Enter the name of the column which contains User IDs to be allowed access to content.
- **Groups column** – Enter the name of the column that contains groups to be allowed access to content.
- **Source URLs column** – Enter the name of the column which contains Source URLs to be indexed.
- **Time stamps column** – Enter the name of the column which contains time stamps. Amazon Q uses time stamp information to detect changes in your content and sync only changed content.
- **Time zones column** – Enter the name of the column which contains time zones for the content to be crawled.
- **Time stamps format** – Enter the name of the column which contains time stamp formats to use to detect content changes and re-sync your content.

14. In **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.

- **Full sync** – Sync all content regardless of the previous sync status.
- **New or modified content sync** – Sync only new and modified documents.
- **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.

For more details, see [Sync mode](#).

15. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
16. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
17. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:
 - a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
 - b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

 **Note**

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

18. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

19. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

 **Note**

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to IBM DB2 using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the configuration parameter to provide a JSON schema with all other configuration information specific to your data source connector.

IBM DB2 JSON schema

The following is the IBM DB2 JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
      "type": "object",
      "properties": {
        "repositoryEndpointMetadata": {
          "type": "object",
          "properties": {
            "dbType": {
              "type": "string",
              "enum": [
                "mysql",
                "db2",
                "postgresql",
                "oracle",
                "sqlserver"
              ]
            },
            "dbHost": {
              "type": "string"
            },
            "dbPort": {
              "type": "string"
            },
            "dbInstance": {
              "type": "string"
            }
          }
        },
        "required": [
          "dbType",
          "dbHost",
```

```
        "dbPort",
        "dbInstance"
    ]
}
},
"required": [
    "repositoryEndpointMetadata"
]
},
"repositoryConfigurations": {
    "type": "object",
    "properties": {
        "document": {
            "type": "object",
            "properties": {
                "fieldMappings": {
                    "type": "array",
                    "items": [
                        {
                            "type": "object",
                            "properties": {
                                "indexFieldName": {
                                    "type": "string"
                                },
                                "indexFieldType": {
                                    "type": "string"
                                },
                                "dataSourceFieldName": {
                                    "type": "string"
                                }
                            }
                        },
                        {
                            "required": [
                                "indexFieldName",
                                "indexFieldType",
                                "dataSourceFieldName"
                            ]
                        }
                    ]
                }
            }
        },
        "required": [
            "fieldMappings"
        ]
    }
}
```

```
    },
    "required": [
    ]
  },
  "additionalProperties": {
    "type": "object",
    "properties": {
      "primaryKey": {
        "type": "string"
      },
      "titleColumn": {
        "type": "string"
      },
      "bodyColumn": {
        "type": "string"
      },
      "sqlQuery": {
        "type": "string",
        "not": {
          "pattern": ";+"
        }
      },
      "timestampColumn": {
        "type": "string"
      },
      "timestampFormat": {
        "type": "string"
      },
      "timezone": {
        "type": "string"
      },
      "changeDetectingColumns": {
        "type": "array",
        "items": {
          "type": "string"
        }
      },
      "allowedUsersColumn": {
        "type": "string"
      },
      "allowedGroupsColumn": {
        "type": "string"
      },
      "sourceURIColumn": {
```

```

        "type": "string"
    },
    "serverlessAurora": {
        "type": "string",
        "enum": ["true", "false"]
    }
},
"required": ["primaryKey", "titleColumn", "bodyColumn", "sqlQuery"]
},
"type" : {
    "type" : "string",
    "pattern": "JDBC"
},
"syncMode": {
    "type": "string",
    "enum": [
        "FORCED_FULL_CRAWL",
        "FULL_CRAWL"
    ]
},
"secretArn": {
    "type": "string"
}
},
"version": {
    "type": "string",
    "anyOf": [
        {
            "pattern": "1.0.0"
        }
    ]
}
},
"required": [
    "connectionConfiguration",
    "repositoryConfigurations",
    "syncMode",
    "additionalProperties",
    "secretArn",
    "type"
]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.
repositoryEndpointMetadata	<p>Required configuration information for connecting your data source.</p> <ul style="list-style-type: none"> • dbType—The type of Java database you are using, whether <code>mysql</code>, <code>db2</code>, <code>postgresql</code>, <code>oracle</code>, or <code>sqlserver</code>. • dbHost—The database host name. • dbPort—The database port. • dbInstance—The database instance.
repositoryConfigurations	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings. Specify the type of data source and the secret ARN.
document	A list of objects that map the attributes or field names of your database content to Amazon Q index field names. For more information, see Mapping data source fields .
additionalProperties	Additional configuration options for your content in your data source. Use to include or exclude specific content in your database data source.
primaryKey	Provide the primary key for the database table. This identifies a table within your database.
titleColumn	Provide the name of the document title column within your database table.

Configuration	Description
bodyColumn	Provide the name of the document title column within your database table.
sqlQuery	Enter SQL query statements like SELECT and JOIN operations. SQL queries must be less than 1000 characters and not contain any semi-colons (;). Amazon Q will crawl all database content that matches your query.
timestampColumn	Enter the name of the column which contains time stamps. Amazon Q uses time stamp information to detect changes in your content and sync only changed content.
timestampFormat	Enter the name of the column which contains time stamp formats to use to detect content changes and re-sync your content.
timezone	Enter the name of the column which contains time zones for the content to be crawled.
changeDetectingColumns	Enter the names of the columns that Amazon Q will use to detect content changes. Amazon Q will re-index content when there is a change in any of these columns
allowedUsersColumns	Enter the name of the column which contains User IDs to be allowed access to content.
allowedGroupsColumn	Enter the name of the column which contains User IDs to be allowed access to content.
sourceURIColumn	Enter the name of the column which contains Source URLs to be indexed.
isSslEnabled	true to add a path to an SSL certificate file stored in an Amazon S3 bucket.

Configuration	Description
type	The type of data source. Specify JDBC as your data source type.
syncMode	<p>Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose</p> <ul style="list-style-type: none">• <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index• <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index• <code>CHANGE_LOG</code> to incrementally crawl only new and modified content each time your data source syncs with your index.
secretArn	<p>The Amazon Resource Name (ARN) of a Secrets Manager secret that contains user name and password required to connect to your database. The secret must contain a JSON structure with the following keys:</p> <pre data-bbox="829 1373 1507 1570">{ "user name": "<i>database user name</i>", "password": "<i>password</i>" }</pre>
version	The version of the template that is currently supported.

How Amazon Q Business connector crawls IBM DB2 ACLs

When you connect a database data source to Amazon Q, Amazon Q crawls user and group information from a column in the source table. You specify this column in the console or using the configuration parameter as part of the `CreateDataSource` operation.

If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

A database data source has the following limitations:

- You can only specify an allow list for a database data source. You can't specify a deny list.
- You can only specify groups. You can't specify individual users for the allow list.
- The database column should be a string containing a semicolon delimited list of groups.

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business IBM DB2 data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a attribute mapping already available, or if you want to map additional

document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q PostgreSQL connector supports the following field mappings:

Supported field mappings

- [Document](#)

Document

JDBC field name	Index field name	Description	Data type
jd_document_id	jd_document_id	Custom	String
jd_document_title	jd_document_title	Custom	String
jd_source_uri	_source_uri	Default	String

IAM role for Amazon Q Business IBM DB2 connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the BatchPutDocument and BatchDeleteDocument operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- Permission to access the SSL certificate stored in your Amazon S3 bucket.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AllowsAmazonQToGetS3Objects",
    "Action": [
      "s3:GetObject"
    ],
    "Resource": [
      "arn:aws:s3:::{{input_bucket_name}}/*"
    ],
    "Effect": "Allow",
    "Condition": {
      "StringEquals": {
        "aws:ResourceAccount": "{{account_id}}"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToGetSecret",
    "Effect": "Allow",
    "Action": [
      "secretsmanager:GetSecretValue"
    ],
    "Resource": [
      "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToDecryptSecret",
    "Effect": "Allow",
    "Action": [
```

```

        "kms:Decrypt"
    ],
    "Resource": [
        "arn:aws:kms:{{region}}:{{account_id}}:key/[key_id]"
    ],
    "Condition": {
        "StringLike": {
            "kms:ViaService": [
                "secretsmanager.*.amazonaws.com"
            ]
        }
    }
},
{
    "Sid": "AllowsAmazonQToIngestDocuments",
    "Effect": "Allow",
    "Action": [
        "qbusiness:BatchPutDocument",
        "qbusiness:BatchDeleteDocument"
    ],
    "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
},
{
    "Sid": "AllowsAmazonQToIngestPrincipalMapping",
    "Effect": "Allow",
    "Action": [
        "qbusiness:PutGroup",
        "qbusiness:CreateUser",
        "qbusiness>DeleteGroup",
        "qbusiness:UpdateUser",
        "qbusiness:ListGroup"
    ],
    "Resource": [
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}",
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}",
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}/data-source/*"
    ]
},
{
    "Sid": "AllowsAmazonQToCreateAndDeleteNI",

```

```

    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2>DeleteNetworkInterface"
    ],
    "Resource": [
      "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
      "arn:aws:ec2:{{region}}:{{account_id}}:security-group/
[[security_group]]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2>DeleteNetworkInterface"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:RequestTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
      },
      "ForAllValues:StringEquals": {
        "aws:TagKeys": [
          "AMAZON_Q"
        ]
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateTags",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateTags"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringEquals": {
        "ec2:CreateAction": "CreateNetworkInterface"
      }
    }
  }
},

```

```

    {
      "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
      "Effect": "Allow",
      "Action": [
        "ec2:CreateNetworkInterfacePermission"
      ],
      "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
      "Condition": {
        "StringLike": {
          "aws:ResourceTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
        }
      }
    },
    {
      "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
      "Effect": "Allow",
      "Action": [
        "ec2:DescribeNetworkInterfaces",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeNetworkInterfaceAttribute",
        "ec2:DescribeVpcs",
        "ec2:DescribeRegions",
        "ec2:DescribeNetworkInterfacePermissions",
        "ec2:DescribeSubnets"
      ],
      "Resource": "*"
    }
  ]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToAssumeRoleForServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
    }
  ]
}

```

```
    "Condition": {
      "StringEquals": {
        "aws:SourceAccount": "{{source_account}}"
      },
      "ArnLike": {
        "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
      }
    }
  }
]
```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Known limitations for the Amazon Q Business IBM DB2 connector

- Deleted database rows will not be tracked in when Amazon Q checks for updated content.
- The size of field names and values in a row of your database can't exceed 400KB.
- If you have a large amount of data in your database data source, and do not want Amazon Q to index all your database content after the first sync, you can choose to sync only new, modified, or deleted documents.

Connecting Jira to Amazon Q Business

Jira is a project management tool for software development, product management, and bug tracking. You can connect your Jira instance to Amazon Q Business—using either the AWS Management Console, CLI, or the [CreateDataSource](#) API—and create an Amazon Q web experience.

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).
- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [Jira connector overview](#)
- [Prerequisites for connecting Amazon Q Business to Jira](#)
- [Setting up Jira for connecting to Amazon Q](#)
- [Connecting Amazon Q Business to Jira using the console](#)
- [Connecting Amazon Q Business to Jira using APIs](#)
- [How Amazon Q Business connector crawls Jira ACLs](#)
- [Amazon Q Business Jira data source connector field mappings](#)
- [IAM role for Amazon Q Business Jira connector](#)
- [Known limitations for the Amazon Q Jira connector](#)
- [Troubleshooting your Amazon Q Business Jira connector](#)

Jira connector overview

The following table gives an overview of the Amazon Q Business Jira connector and its supported features.

Category	Feature	Support
Security	Authentication type	Basic
	Authentication credentials	<ul style="list-style-type: none"> • Jira ID • Jira password/token
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Identity crawling	No
	VPC	Yes
Crawl features	Custom objects	Yes
	Custom metadata	Yes
	Entities	Yes. The following entities are supported: <ul style="list-style-type: none"> • Projects

Category	Feature	Support
		<ul style="list-style-type: none"> • Issues • Comments • Attachments • Worklogs
	Field mappings	Yes. Supports both default and custom field mappings. For more information, see Field mappings .
	Filters	<p>Yes. The following filters are supported:</p> <ul style="list-style-type: none"> • Include specific projects • Include/exclude statuses • Include/exclude comments • Include/exclude attachments • Include/exclude worklogs • Include/exclude bugs • Include/exclude epic • Include/exclude story • Include/exclude task • Include/exclude by file name • Include/exclude by file type • Include/exclude by file path
	Sync mode	Supports full and incremental sync.
	File types	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to Jira

Before you begin, make sure that you have completed the following prerequisites.

In Jira, make sure you have:

- Created Jira API token authentication credentials that include a Jira ID (email ID with domain) and a Jira credential (Jira API token). See [Atlassian documentation on managing API tokens](#).
- Noted the Jira account URL from your Jira account settings. For example, *https://company.atlassian.net/*.
- Noted your Jira project key ID from your Jira project settings if you want to crawl only specific Jira projects.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your Jira authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Setting up Jira for connecting to Amazon Q

Before you connect Jira to Amazon Q, you need to create and retrieve the Jira credentials you will use to connect Jira to Amazon Q.

The following procedures gives you an overview of how to configure Jira for connecting with Amazon Q.

Topics

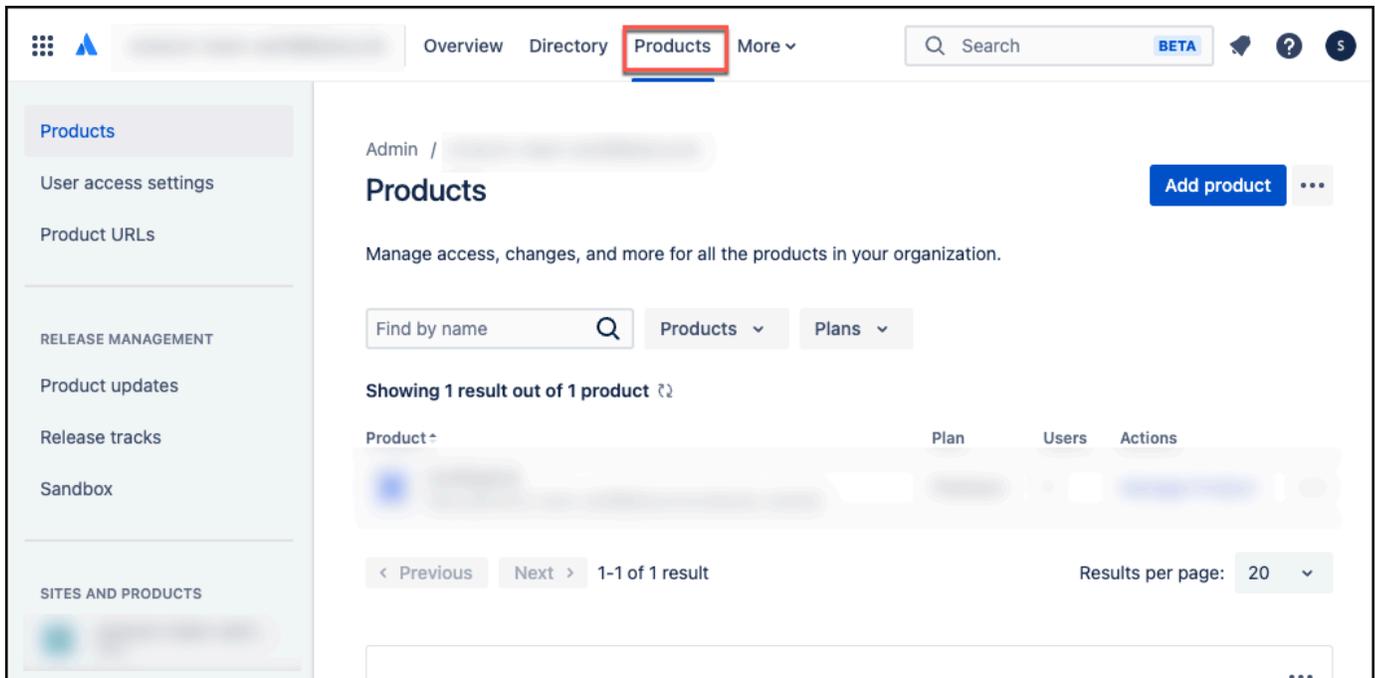
- [Configuring basic authentication](#)
- [Retrieving project key](#)

Configuring basic authentication

You can connect Amazon Q to Jira using basic authentication credentials. The following procedure gives you an overview of how to configure Jira to connect to Amazon Q using basic authentication.

Configuring Jira basic authentication for Amazon Q

1. Sign up for an Atlassian account from <https://atlassian.com/>. Note the email id, including domain, that you logged in with. You will input this later as the Jira ID when you connect to Amazon Q.
2. Navigate to Atlassian account from <https://admin.atlassian.com/>. This is where you will configure your Jira instance.
3. From the top navigation menu, select **Products**. Then, select **Add product**.



4. On the **Select product** page, select **Jira Software**. Then, select **Select**.

Select product  Add product

Select a product

Select a product to add to your organization. Try it out for free before you subscribe to a plan that works for you.

 **Atlassian Access** Security controls for users, data and devices DETAILS ▾

 **Bitbucket** Git code management DETAILS ▾

 **Confluence** Document collaboration DETAILS ▾

 **Jira Product Discovery** Dynamic product discovery DETAILS ▾

  **Jira Service Management** High-velocity ITSM DETAILS ▾

 **Jira Software** Project and issue tracking DETAILS ▾

 **Jira Work Management**

5. On the **Add product** page, select **Create new site**. Then, for **Site name**, add a name for your Jira site. Copy the site name, including the domain name. For example: *https://company.atlassian.net/*. You will input this as your **Jira Account URL** when you connect to Amazon Q.

Select **Agree and add**.

Select product

Add product

1

Jira Software

Start your 7-day free Cloud Standard trial

Add Jira Software to a new site or to one of your existing sites. [Learn more about adding products](#)

Create new site

Site name 

2

 .atlassian.net 

The site name must be at least three characters (numerals or lowercase letters only) and can't start or end with a hyphen (-).

Add to an existing site

This site is protected by reCAPTCHA and the Google [Privacy Policy](#) and [Terms of Service](#) apply. By clicking below, you agree to the Atlassian [Cloud Terms of Service](#) and [Privacy Policy](#).

Back

Agree and add

3

6. Log in to your Atlassian account from <https://atlassian.com/>.
7. From the top navigation menu, navigate to **Security**. Then, from **API Tokens**, select **Create and manage API tokens**.

Security 1

Change your password

When you change your password, we keep you logged in to this device but may log you out from your other devices.

Current password *

New password *

Save changes

Two-step verification

Keep your account extra secure with a second login step. [Learn more](#)

[Manage two-step verification](#)

API tokens

A script or other process can use an API token to perform basic authentication with Jira Cloud applications or Confluence Cloud. You must use an API token if the Atlassian account you authenticate with has had two-step verification enabled. You should treat API tokens as securely as any other password. [Learn more](#)

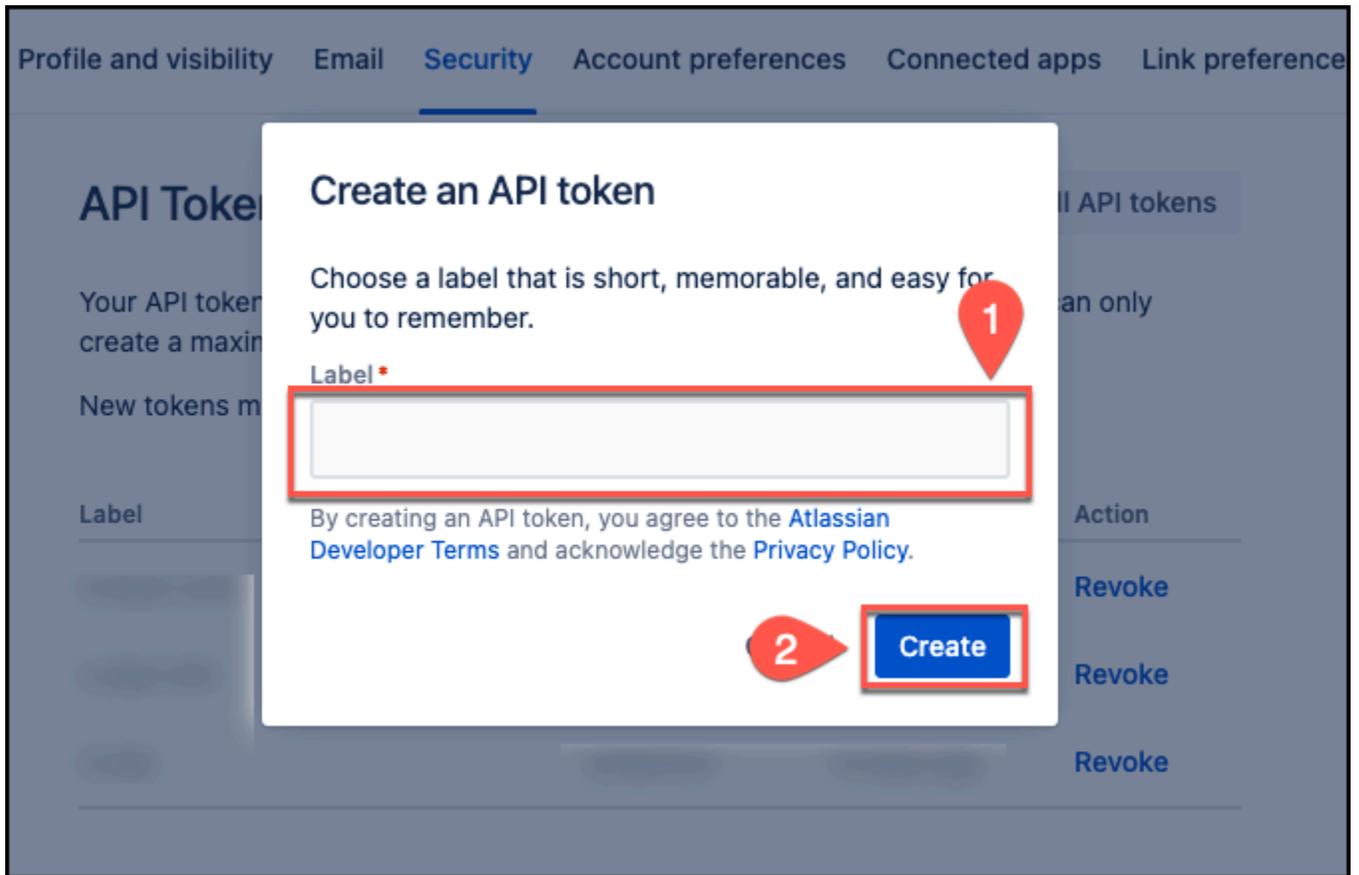
[Create and manage API tokens](#) 2

Recent devices

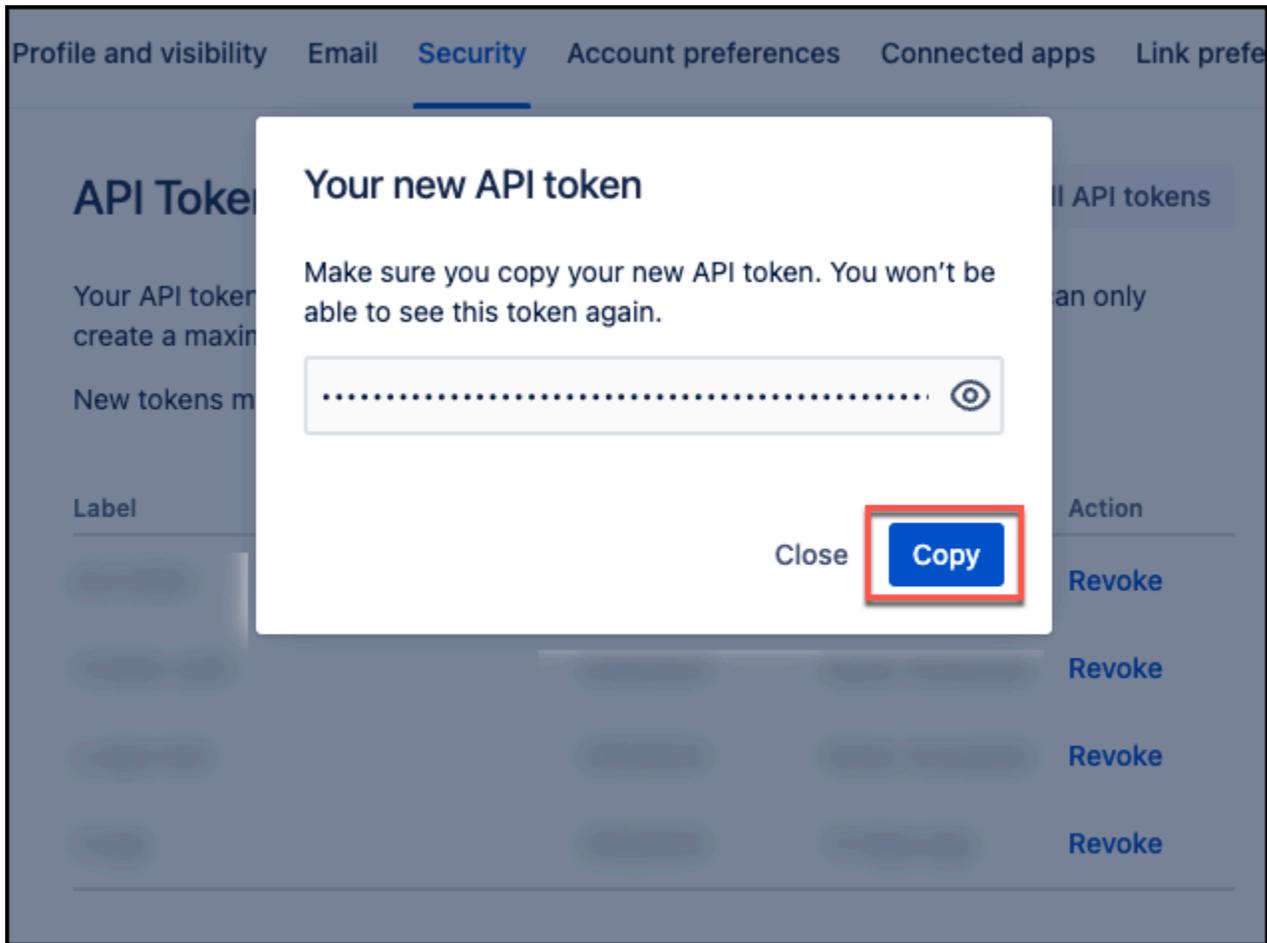
If you've lost one of your devices or notice any suspicious activity, log out of all your devices and take steps to secure your account. [Learn more](#)

[View and manage recent devices](#)

8. In **API Tokens**, for **Create an API token**, in **Label**, add a label name. Then, select **Create**.



9. From the **Your new API token** dialog box, copy the API token and save it in a text editor of your choice. You can't retrieve the API token once you close the dialog box. You use the API token to connect Jira to Amazon Q.



You now have the username, Jira URL, and Jira API token you need to connect to Amazon Q with basic authentication.

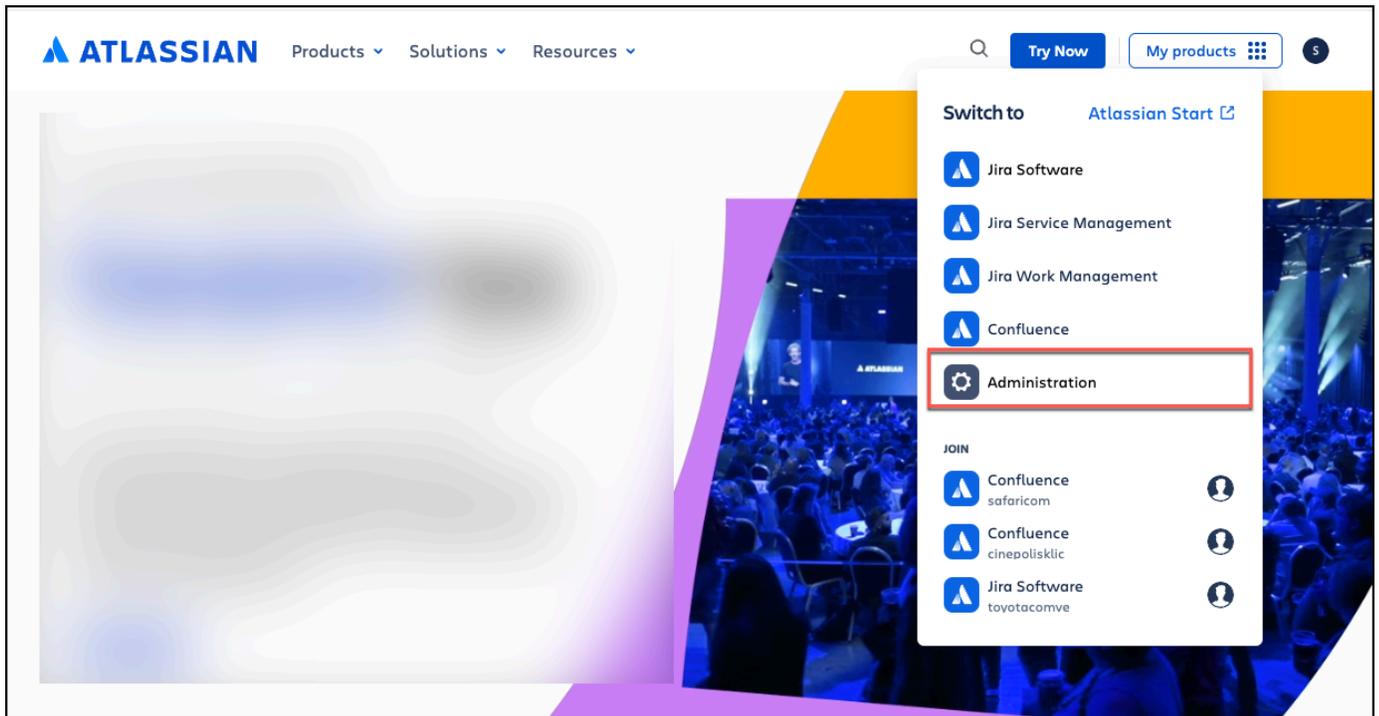
For more information, see [Manage API tokens for your Atlassian account](#) in Atlassian Support.

Retrieving project key

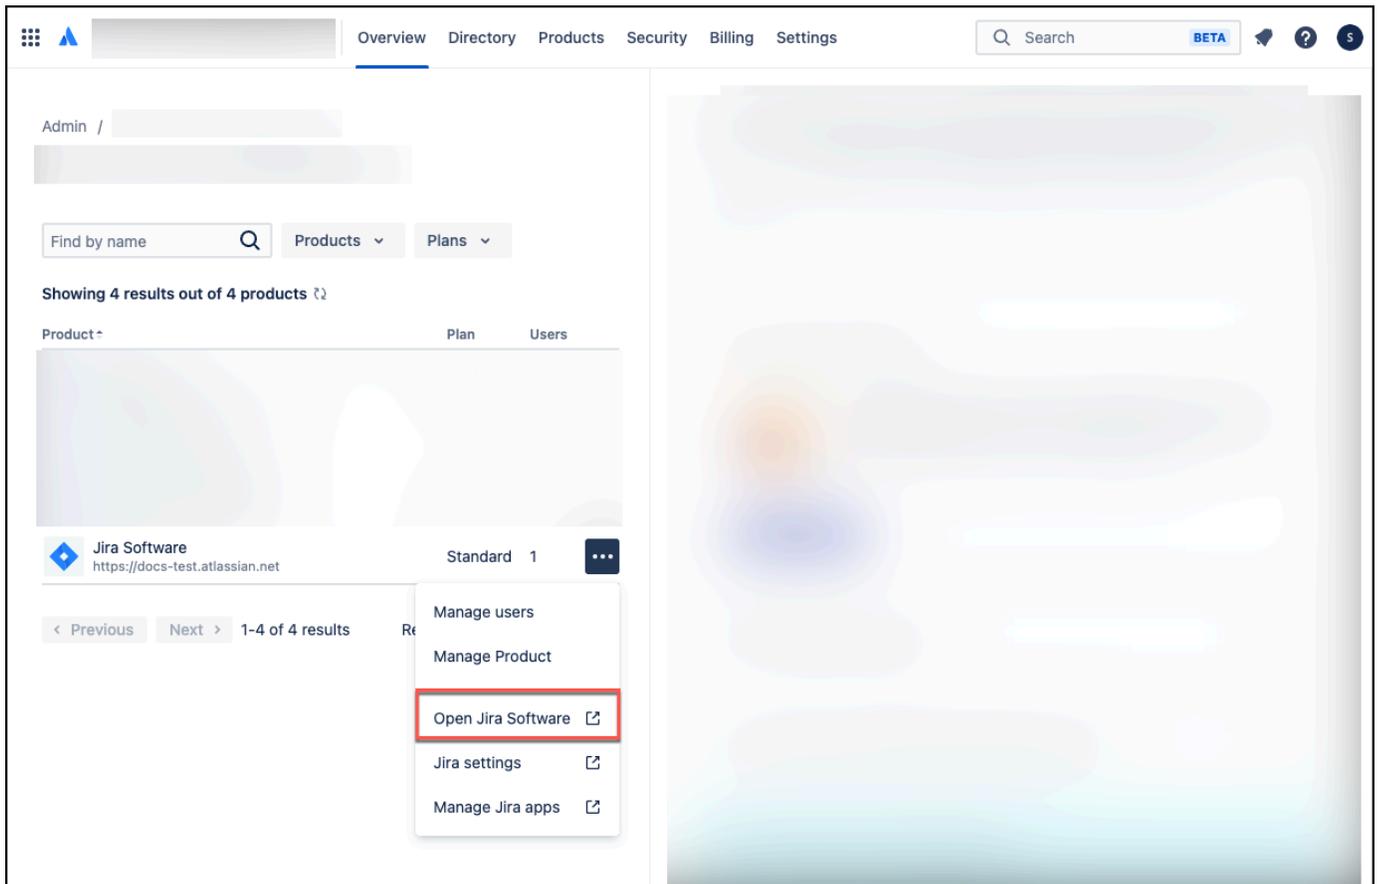
Amazon Q connector gives you the ability to crawl specific Jira projects instead of crawling all Jira projects. To crawl a specific Jira project, you need to retrieve its **Project Key**. Then, when you connect Jira to Amazon Q, you provide the specific project key you want to crawl in the **Sync scope** section. The following procedure gives you an overview of how to retrieve a Jira **Project Key**.

Retrieving a Jira project key

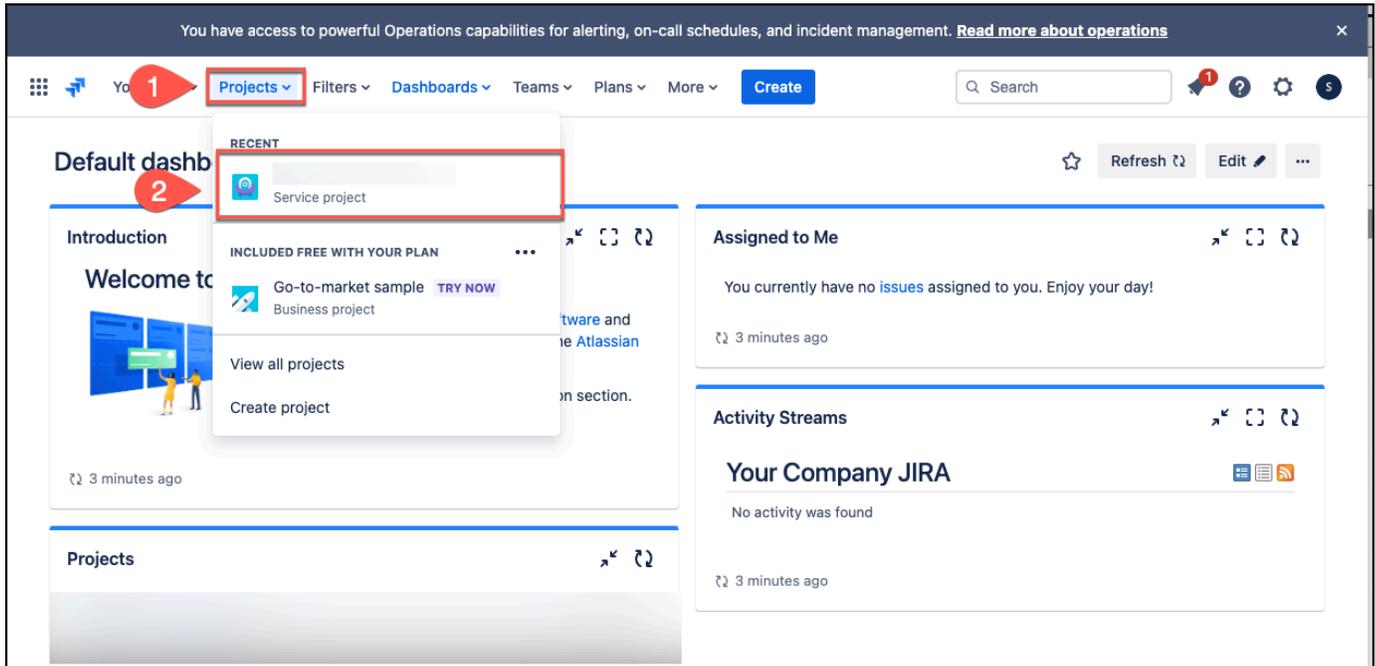
1. Log in to your Atlassian account from <https://atlassian.com>.
2. From the profile top-right navigation menu, choose **My products**. Then choose **Administration**.



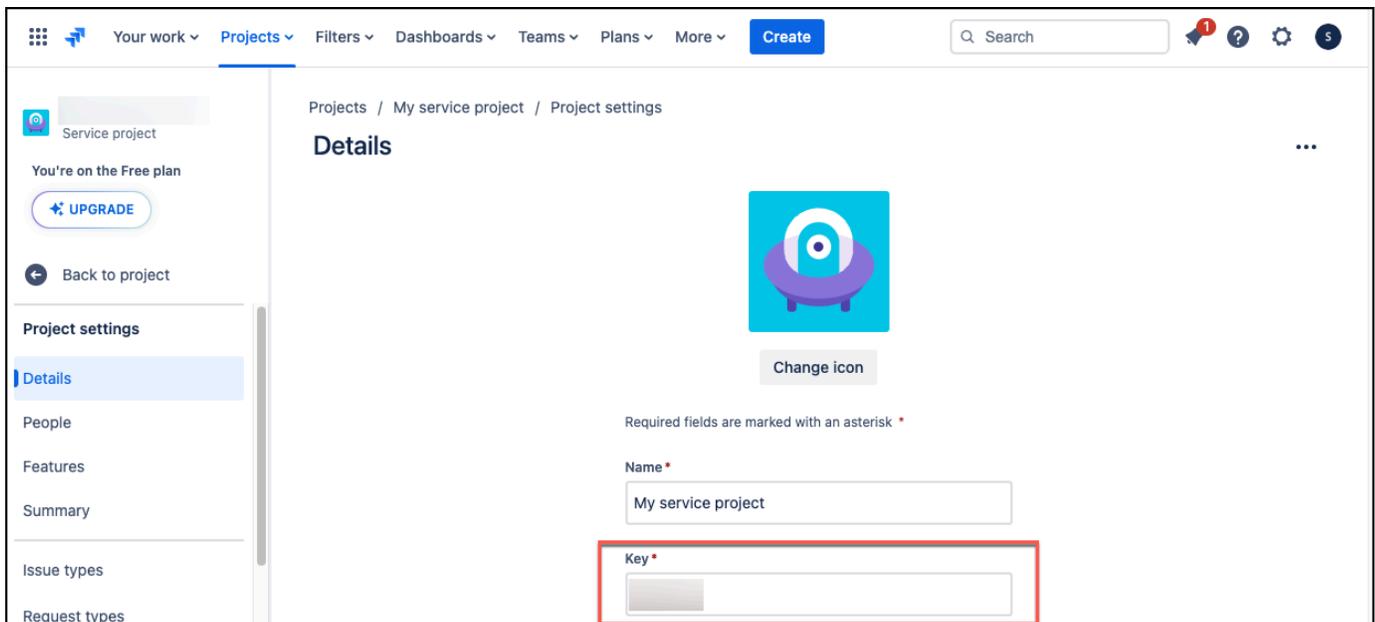
- From the admin overview page, navigate to your **Jira Software** instance, and then open the settings menu. Select **Open Jira Software**.



- From the **Default dashboard page**, from the top navigation menu choose **Projects**, and then select your project.
- On your project page, from the left navigation menu choose **Project settings**.



- The **Details** page will display your project key under **Key**.



Connecting Amazon Q Business to Jira using the console

The following procedure outlines how to connect Amazon Q Business to Jira using the AWS Management Console.

Connecting Amazon Q to Jira

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **Jira** page, enter the following information:
6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. **Source** – Enter your **Jira URL**. For example: *https://company.atlassian.net/*.
8. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

Note

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

9. **Authentication** – Enter the following information for your **AWS Secrets Manager secret**.
 - a. **Secret name** – A name for your secret.
 - b. **Jira ID** – Your Jira email id, with domain.
 - c. **Password/Token** – Your Jira API token.
10. **Configure VPC and security group** – *optional* – Choose whether you want to use a VPC. If you do, enter the following information:

- a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
- b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

11. **Identity crawler** – Choose to activate Amazon Q identity crawler to sync identity information. For more information, see [Identity crawler](#).
12. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

13. In **Sync scope**, enter the following information:
 - a. Choose to either sync **All projects** or sync **Only specific projects**. If you choose to sync **Only specific projects**, enter the **Jira Project Key ID**.
 - b. **Additional configuration – optional** – Choose from the following options to limit the scope for content to be indexed. Otherwise, all content will be synced by default.
 - **Statuses** – Add status values to index.
 - **Additional elements** – Choose whether to index **Comments**, **Attachments**, or **Worklogs**.
 - **Issue types** – Choose the issues types you want to index.
 - **Regex patterns** – Add regex patterns to include or exclude file names, file types, or file paths. You can have a total of 100 patterns.
14. For **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.
 - **Full sync** – Sync all content regardless of the previous sync status.
 - **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.
15. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).

16. **Tags - *optional*** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
17. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:
 - a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
 - b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

 **Note**

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

18. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

19. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

 **Note**

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to Jira using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the configuration parameter to provide a JSON schema with all other configuration information specific to your data source connector.

For an example of the API request, see [CreateDataSource](#) in the Amazon Q API Reference.

Jira JSON schema

The following is the Jira JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
      "type": "object",
      "properties": {
        "repositoryEndpointMetadata": {
          "type": "object",
          "properties": {
            "jiraAccountUrl": {
              "type": "string",
              "pattern": "https://.*"
            }
          },
          "required": [
            "jiraAccountUrl"
          ]
        }
      },
      "required": [
        "repositoryEndpointMetadata"
      ]
    },
    "repositoryConfigurations": {
      "type": "object",
      "properties": {
        "attachment": {
          "type": "object",
          "properties": {
            "fieldMappings": {
              "type": "array",
              "items": [
                {
                  "type": "object",
```



```

    },
    "indexFieldType": {
      "type": "string",
      "enum": [
        "STRING",
        "STRING_LIST",
        "DATE"
      ]
    },
    "dataSourceFieldName": {
      "type": "string"
    },
    "dateFieldFormat": {
      "type": "string",
      "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
    }
  },
  "required": [
    "indexFieldName",
    "indexFieldType",
    "dataSourceFieldName"
  ]
}
]
}
},
"required": [
  "fieldMappings"
]
},
"issue": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": [
        {
          "type": "object",
          "properties": {
            "indexFieldName": {
              "type": "string"
            },
            "indexFieldType": {
              "type": "string",

```

```

        "enum": [
            "STRING",
            "STRING_LIST",
            "DATE"
        ]
    },
    "dataSourceFieldName": {
        "type": "string"
    },
    "dateFieldFormat": {
        "type": "string",
        "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
    }
},
"required": [
    "indexFieldName",
    "indexFieldType",
    "dataSourceFieldName"
]
}
]
}
},
"required": [
    "fieldMappings"
]
},
"project": {
    "type": "object",
    "properties": {
        "fieldMappings": {
            "type": "array",
            "items": [
                {
                    "type": "object",
                    "properties": {
                        "indexFieldName": {
                            "type": "string"
                        },
                    },
                    "indexFieldType": {
                        "type": "string",
                        "enum": [
                            "STRING",
                            "STRING_LIST",

```

```

        "DATE"
      ]
    },
    "dataSourceFieldName": {
      "type": "string"
    },
    "dateFieldFormat": {
      "type": "string",
      "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
    }
  },
  "required": [
    "indexFieldName",
    "indexFieldType",
    "dataSourceFieldName"
  ]
}
]
}
},
"required": [
  "fieldMappings"
]
},
"worklog": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": [
        {
          "type": "object",
          "properties": {
            "indexFieldName": {
              "type": "string"
            },
            "indexFieldType": {
              "type": "string",
              "enum": [
                "STRING",
                "STRING_LIST",
                "DATE"
              ]
            }
          }
        }
      ]
    }
  }
},

```

```

        "dataSourceFieldName": {
            "type": "string"
        },
        "dateFieldFormat": {
            "type": "string",
            "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
        }
    },
    "required": [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
    ]
}
]
}
},
"required": [
    "fieldMappings"
]
}
},
"additionalProperties": {
    "type": "object",
    "properties": {
        "isCrawlAcl": {
            "type": "boolean"
        },
        "fieldForUserId": {
            "type": "string"
        },
        "issuetype": {
            "type": "array",
            "items": {
                "type": "string",
                "enum": [
                    "Bug",
                    "Story",
                    "Epic",
                    "Task"
                ]
            }
        }
    }
},

```

```
    "status": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    "project": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    "issueSubEntityFilter": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    "inclusionPatterns": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    "exclusionPatterns": {
      "type": "array",
      "items": {
        "type": "string"
      }
    }
  },
  "required": []
},
"type": {
  "type": "string",
  "pattern": "JIRA"
},
"enableIdentityCrawler": {
  "type": "boolean"
},
"syncMode": {
  "type": "string",
  "enum": [
    "FULL_CRAWL",
```

```

        "FORCED_FULL_CRAWL",
        "CHANGE_LOG"
    ]
},
"secretArn": {
    "type": "string",
    "minLength": 20,
    "maxLength": 2048
}
},
"version": {
    "type": "string",
    "anyOf": [
        {
            "pattern": "1.0.0"
        }
    ]
}
},
"required": [
    "connectionConfiguration",
    "repositoryConfigurations",
    "syncMode",
    "additionalProperties",
    "secretArn",
    "type",
    "enableIdentityCrawler"
]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.
repositoryEndpointMetadata	The endpoint information for the data source.
jiraAccountUrl	Enter the Jira account URL from your Jira account settings. For example, <i>https://company.atlassian.net/</i> .

Configuration	Description
<code>repositoryConfigurations</code>	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings.
<ul style="list-style-type: none"> • <code>attachment</code> • <code>comment</code> • <code>issue</code> • <code>project</code> • <code>worklog</code> 	A list of objects that map the attributes or field names of your Jira pages and assets to Amazon Q index field names.
<code>additionalProperties</code>	Additional configuration options for your content in your data source.
<code>isCrawlAcl</code>	Specify <code>true</code> to crawl access control information from documents.
<code>fieldForUserId</code>	Specify field to use for <code>UserId</code> for ACL crawling.
<ul style="list-style-type: none"> • <code>issuetype</code> • <code>status</code> • <code>project</code> • <code>issueSubEntityFilter</code> 	Choose to customize the scope of your crawl with specific entities. You can add specific status types, additional elements, and issue types to crawl.
<ul style="list-style-type: none"> • <code>inclusionPatterns</code> 	A list of regular expression patterns to include specific content in your Jira data source. Content that matches the patterns are included in the index. Contents that doesn't match the pattern are excluded from the index. If any content matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence, and the content isn't included in the index.

Configuration	Description
<ul style="list-style-type: none"> <code>exclusionPatterns</code> 	<p>A list of regular expression patterns to exclude specific content in your Jira data source. Content that matches the patterns are excluded from the index. Content that doesn't match the patterns are included in the index. If any content matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence, and the content isn't included in the index.</p>
<code>type</code>	<p>The type of data source. Specify JIRA as your data source type.</p>
<code>enableIdentityCrawler</code>	<p>Specify <code>true</code> to use the Amazon Q identity crawler to sync identity/principal information on users and groups with access to specific documents.</p>
<code>syncMode</code>	<p>Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose between the following options:</p> <ul style="list-style-type: none"> Use <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index. Use <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index. Use <code>CHANGE_LOG</code> to incrementally crawl only new and modified content each time your data source syncs with your index.

Configuration	Description
secretArn	<p>The Amazon Resource Name (ARN) of an AWS Secrets Manager secret that contains the key-value pairs required to connect to your Jira. The secret must contain a JSON structure with the following keys:</p> <pre>{ "Jira ID": "<i>Jira user name or email host URL</i>", "Password/Token": "<i>Jira API token</i>" }</pre>
version	<p>The version of this template that's currently supported.</p>

How Amazon Q Business connector crawls Jira ACLs

When you connect an Jira data source to Amazon Q Business, Amazon Q Business crawls ACL information attached to a document (user and group information) from your Jira instance. If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

The Jira user IDs are mapped as follows:

- `_user_id`—User IDs exist in Jira on files where there are set access permissions. They are mapped from the user emails as the user IDs in Jira.

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business Jira data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q Business enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a attribute mapping already available, or if you want to map additional document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q Jira connector supports the following entities and the associated reserved and custom attributes.

Supported entities and field mappings

- [Projects](#)
- [Issues](#)
- [Comments](#)
- [Attachments](#)

- [Worklogs](#)

Projects

Jira field name	Index field name	Description	Data type
title	j_title	Custom	String
project_key	j_project_key	Custom	String
lead	j_lead	Custom	String list
url	_source_uri	Default	String

Issues

Jira field name	Index field name	Description	Data type
title	j_title	Custom	String
issue_key	j_issue_key	Custom	String
status	j_status	Custom	String
project_name	j_project_name	Custom	String
projectKey	j_project_key	Custom	String
authors	_authors	Default	String list
assignee	j_assignee	Custom	String
created_at	_created_at	Default	Date
updated_at	_last_updated_at	Default	Date
url	_source_uri	Default	String
issue_type	j_issue_type	Custom	String

Jira field name	Index field name	Description	Data type
priority	j_priority	Custom	String
resolution	j_resolution	Custom	String
affects_version	j_affects_version	Custom	String
fix_version	j_fix_version	Custom	String
labels	j_labels	Custom	String
environment	j_environment	Custom	String
reporter	j_reporter	Custom	String
votes	j_votes	Custom	String
watchers	j_watchers	Custom	String
due	j_due	Custom	String
resolved	j_resolved	Custom	String

Comments

Jira field name	Index field name	Description	Data type
authors	_authors	Default	String list
title	j_title	Custom	String
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date
project_name	j_project_name	Custom	String
project_key	j_project_key	Custom	String

Jira field name	Index field name	Description	Data type
issue_key	j_issue_key	Custom	String
url	_source_uri	Default	String

Attachments

Jira field name	Index field name	Description	Data type
title	j_title	Custom	String
authors	_authors	Default	String list
size	j_size	Custom	String
createdAt	_created_at	Default	Date
url	_source_uri	Default	String
project_name	j_project_name	Custom	String
project_key	j_project_key	Custom	String
issue_key	j_issue_key	Custom	String

Worklogs

Jira field name	Index field name	Description	Data type
title	j_title	Custom	String
authors	_authors	Default	String list
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date
url	_source_uri	Default	String

Jira field name	Index field name	Description	Data type
project_name	j_project_name	Custom	String
project_key	j_project_key	Custom	String
issue_key	j_issue_key	Custom	String

IAM role for Amazon Q Business Jira connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the `BatchPutDocument` and `BatchDeleteDocument` operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToGetSecret",
      "Effect": "Allow",
      "Action": [
        "secretsmanager:GetSecretValue"
      ],
      "Resource": [
```

```

    "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
  ]
},
{
  "Sid": "AllowsAmazonQToDecryptSecret",
  "Effect": "Allow",
  "Action": [
    "kms:Decrypt"
  ],
  "Resource": [
    "arn:aws:kms:{{region}}:{{account_id}}:key/[key_id]"
  ],
  "Condition": {
    "StringLike": {
      "kms:ViaService": [
        "secretsmanager.*.amazonaws.com"
      ]
    }
  }
},
{
  "Sid": "AllowsAmazonQToIngestDocuments",
  "Effect": "Allow",
  "Action": [
    "qbusiness:BatchPutDocument",
    "qbusiness:BatchDeleteDocument"
  ],
  "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
},
{
  "Sid": "AllowsAmazonQToIngestPrincipalMapping",
  "Effect": "Allow",
  "Action": [
    "qbusiness:PutGroup",
    "qbusiness:CreateUser",
    "qbusiness>DeleteGroup",
    "qbusiness:UpdateUser",
    "qbusiness:ListGroup"
  ],
  "Resource": [
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}",
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}",

```

```

        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}/data-source/*"
    ]
},
{
    "Sid": "AllowsAmazonQToCreateAndDeleteNI",
    "Effect": "Allow",
    "Action": [
        "ec2:CreateNetworkInterface",
        "ec2:DeleteNetworkInterface"
    ],
    "Resource": [
        "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
        "arn:aws:ec2:{{region}}:{{account_id}}:security-group/[{{security_group}}]"
    ]
},
{
    "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
    "Effect": "Allow",
    "Action": [
        "ec2:CreateNetworkInterface",
        "ec2:DeleteNetworkInterface"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
        "StringLike": {
            "aws:RequestTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
        },
        "ForAllValues:StringEquals": {
            "aws:TagKeys": [
                "AMAZON_Q"
            ]
        }
    }
},
{
    "Sid": "AllowsAmazonQToCreateTags",
    "Effect": "Allow",
    "Action": [
        "ec2:CreateTags"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
        "StringEquals": {

```



```

        "ec2:CreateAction": "CreateNetworkInterface"
    }
}
},
{
    "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
    "Effect": "Allow",
    "Action": [
        "ec2:CreateNetworkInterfacePermission"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
        "StringLike": {
            "aws:ResourceTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
        }
    }
},
{
    "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
    "Effect": "Allow",
    "Action": [
        "ec2:DescribeNetworkInterfaces",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeNetworkInterfaceAttribute",
        "ec2:DescribeVpcs",
        "ec2:DescribeRegions",
        "ec2:DescribeNetworkInterfacePermissions",
        "ec2:DescribeSubnets"
    ],
    "Resource": "*"
}
]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```

{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "AllowsAmazonQServicePrincipal",
            "Effect": "Allow",
            "Principal": {

```

```

    "Service": "qbusiness.amazonaws.com"
  },
  "Action": "sts:AssumeRole",
  "Condition": {
    "StringEquals": {
      "aws:SourceAccount": "{{source_account}}"
    },
    "ArnEquals": {
      "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
    }
  }
}
]
}

```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Known limitations for the Amazon Q Jira connector

The Amazon Q Jira connector has the following known limitations:

- Deleted Issues in Jira are not available through Jira APIs. The Amazon Q Jira connector won't be able to fetch information about deleted Jira issues during incremental syncs.
- Private and Empty projects aren't crawled by the Amazon Q Jira connector.

Troubleshooting your Amazon Q Business Jira connector

The following table provides information about error codes you may see for the Jira connector and suggested troubleshooting actions.

Error code	Error message	
JIRA-5100	There was a problem while retrieving access token. Access token should not be null or empty.	

Error code	Error message	
JIRA-5101	There was an error parsing the field value. The size has exceeded the maximum allowable limit.	
JIRA-5102	Jira inclusion pattern list is too large.	
JIRA-5103	Some of the inclusion objects exceed the character limit.	
JIRA-5104	Jira exclusion pattern size list is too large.	
JIRA-5105	Some of the exclusion objects exceed the character limit.	
JIRA-5106	There was a problem while retrieving refresh token. Refresh token should not be null or empty.	
JIRA-5107	There was a problem while retrieving Jira Credential. Jira Credential should not be null or empty.	
JIRA-5108	There was a problem while retrieving Jira Id. Jira Id should not be null or empty.	
JIRA-5109	There was a problem while retrieving Auth Type. Auth Type should not be null or empty.	
JIRA-5110	There was a problem while retrieving Jira Account Url. Jira Account Url should not be null or empty.	
JIRA-5111	JIRA Issue Sub Entity Filter list size is too large.	

Error code	Error message	
JIRA-5112	Some of the Jira Issue Sub Entity Filter objects exceed the character limit.	
JIRA-5113	Jira Issue Status Filter list size is too large.	
JIRA-5114	Some of the Jira Issue Status Filter objects exceeded the character limit.	
JIRA-5115	Jira Issue Type Filter list size is too large.	
JIRA-5116	Some of the Jira Issue Type Filter objects exceed the character limit.	
JIRA-5117	Jira Project Key Filter list size is too large.	
JIRA-5118	Some of the JIRA Project Key Filter objects exceed the character limit.	
JIRA-5119	Project specific field mappings are not configured for connector.	
JIRA-5120	Issue specific field mappings are not configured for connector.	
JIRA-5121	Comment specific field mappings are not configured for connector.	
JIRA-5122	Attachment specific field mappings are not configured for connector.	
JIRA-5123	Worklog specific field mappings are not configured for connector.	

Error code	Error message
JIRA-5124	There was a problem while retrieving crawl type. Crawl Type should not be null or empty.

Connecting Microsoft Exchange to Amazon Q Business

Microsoft Exchange is an enterprise collaboration tool for messaging, meetings, and file sharing. You can connect Microsoft Exchange instance to Amazon Q Business—using either the AWS Management Console or the [CreateDataSource](#) API—and create an Amazon Q web experience.

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).
- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [Microsoft Exchange connector overview](#)
- [Prerequisites for connecting Amazon Q Business to Microsoft Exchange](#)
- [Connecting Amazon Q Business to Microsoft Exchange using the console](#)
- [Connecting Amazon Q Business to Microsoft Exchange using APIs](#)
- [How Amazon Q Business connector crawls Exchange ACLs](#)
- [Amazon Q Business Microsoft Exchange data source connector field mappings](#)
- [IAM role for Amazon Q Business Microsoft Exchange connector](#)
- [Troubleshooting your Amazon Q Business Microsoft Exchange connector](#)

Microsoft Exchange connector overview

The following table gives an overview of the Amazon Q Business Microsoft Exchange connector and its supported features.

Category	Feature	Support
Security	Authentication type	OAuth 2.0
	Authentication credentials	<ul style="list-style-type: none"> • Microsoft Exchange Client ID • Microsoft Exchange Client secret
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Identity crawling	No
	VPC	Yes
Crawl features	Custom metadata	No
	Entities	<p>Yes. The following entities are supported:</p> <ul style="list-style-type: none"> • Mail • Calendar • Attachment • OneNotes • Contacts
	Field mappings	Yes. Supports both default and custom field mappings. For more information, see Field mappings .
	Filters	<p>Yes. The following filters are supported:</p> <ul style="list-style-type: none"> • Include/exclude Calendars • Include/exclude OneNotes • Include/exclude Contacts • Include/exclude using file user email ID • Include/exclude using date • Include/exclude using email to, from, subjects, domains

Category	Feature	Support
		<ul style="list-style-type: none"> • Include/exclude by file name regex patterns • Include/exclude by file type regex patterns
	Sync mode	Supports full and incremental sync.
	File types	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to Microsoft Exchange

In Microsoft Exchange, make sure you have:

- Created a Microsoft Exchange account in Office 365.
- Copied your Microsoft 365 tenant ID. You can find your tenant ID in the **Properties** of your Azure Active Directory Portal. For more information, see [Find your Microsoft 365 tenant ID](#) on the Microsoft website.
- Configured an OAuth 2.0 credential token containing a client ID and client secret.
- Added the following permissions for the connector application:

Microsoft Graph	Office 365 Exchange Online
<ul style="list-style-type: none"> • Mail.Read (Application) • Mail.ReadBasic (Application) • Mail.ReadBasic.All (Application) • Calendars.Read (Application) • User.Read.All (Application) • Contacts.Read (Application) • Notes.Read.All (Application) • Directory.Read.All (Application) • EWS.AccessAsUser.All (Delegated) 	<ul style="list-style-type: none"> • full_access_as_app (Application)

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your Microsoft Exchange authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Connecting Amazon Q Business to Microsoft Exchange using the console

The following procedure outlines how to connect Amazon Q Business to Microsoft Exchange using the AWS Management Console.

Connecting Amazon Q to Microsoft Exchange

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **Microsoft Exchange** page, enter the following information:
6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. In **Source**, enter the following information:
 - **Tenant ID** – Enter your tenant id. Your Microsoft tenant ID is a globally unique identifier that's necessary to configure each connector instance. Your tenant ID is different from your

organization name or domain and can be found in the properties section of your Microsoft account dashboard.

8. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

 **Note**

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

9. **Authentication** – Choose between **New** and **Existing**.

- If you choose **Existing**, select an existing secret for **Select secret**.

If you choose **New**, enter the following information in the **New AWS Secrets Manager secret** section:

- i. **Secret name** – A name for your secret.
- ii. For **Client ID, Client secret** – Enter the authentication credential values that you generated from your Exchange account.

10. **Configure VPC and security group** – *optional* – Choose whether you want to use a VPC. If you do, enter the following information:

- a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
- b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

11. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

12. In **Sync scope**, choose from the following options:

- **UserIDs** – Select to filter content by specific user email IDs.
 - **User email ID** – Upload a file with user email ids to filter content by. Email IDs must be formatted on a separate line in the file.
13. For **Additional configuration – optional**, choose from the following options:
- **Entity types** – Choose whether you want to crawl the following entities: **Calendar**, **OneNotes**, and **Contacts**.
 - **Calendar crawling** – Enter the date range for which the connector will crawl your calendar content.
 - **Include email** – Enter the email from domains, email to domains, and subjects you wish to include or exclude in your application.
 - **Shared folders access** – Enable ACL crawling for shared folders.
 - **Regex for domains** – Add patterns to include and exclude certain email domains from your application.
 - **Regex patterns** – Add regular expression patterns to include or exclude certain files. You can add up to 100 patterns.
14. In **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.
- **Full sync** – Sync all content regardless of the previous sync status.
 - **New or modified content sync** – Sync only new and modified documents.
 - **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.
- For more details, see [Sync mode](#).
15. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
16. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
17. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:

- a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
- b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

 **Note**

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

18. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

19. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

 **Note**

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to Microsoft Exchange using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the `configuration` parameter to provide a JSON schema with all other configuration information specific to your data source connector.

For an example of the API request, see [CreateDataSource](#) in the Amazon Q API Reference.

Microsoft Exchange JSON schema

The following is the Microsoft Exchange JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
      "type": "object",
      "properties": {
        "repositoryEndpointMetadata": {
          "type": "object",
          "properties": {
            "tenantId": {
              "type": "string",
              "pattern": "^[0-9a-f]{8}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]
{12}$",
              "minLength": 36,
              "maxLength": 36
            }
          },
          "required": ["tenantId"]
        }
      }
    },
    "repositoryConfigurations": {
      "type": "object",
      "properties": {
        "email": {
          "type": "object",
          "properties": {
            "fieldMappings": {
              "type": "array",
              "items": [
                {
                  "type": "object",
                  "properties": {
                    "indexFieldName": {
                      "type": "string"
                    },
                    "indexFieldType": {
                      "type": "string",
                      "enum": ["STRING", "STRING_LIST", "DATE"]
                    }
                  }
                }
              ]
            }
          }
        }
      }
    }
  }
}
```

```

        },
        "dataSourceFieldName": {
            "type": "string"
        },
        "dateFieldFormat": {
            "type": "string",
            "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
        }
    },
    "required": [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
    ]
}
]
}
},
"required": [
    "fieldMappings"
]
},
"attachment": {
    "type": "object",
    "properties": {
        "fieldMappings": {
            "type": "array",
            "items": [
                {
                    "type": "object",
                    "properties": {
                        "indexFieldName": {
                            "type": "string"
                        },
                        "indexFieldType": {
                            "type": "string",
                            "enum": ["STRING", "DATE", "LONG"]
                        },
                        "dataSourceFieldName": {
                            "type": "string"
                        },
                        "dateFieldFormat": {
                            "type": "string",
                            "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
                        }
                    }
                }
            ]
        }
    }
}

```

```

        }
      },
      "required": [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
      ]
    }
  ]
}
},
"required": [
  "fieldMappings"
]
},
"calendar": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": [
        {
          "type": "object",
          "properties": {
            "indexFieldName": {
              "type": "string"
            },
            "indexFieldType": {
              "type": "string",
              "enum": ["STRING", "STRING_LIST", "DATE"]
            },
            "dataSourceFieldName": {
              "type": "string"
            },
            "dateFieldFormat": {
              "type": "string",
              "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
            }
          }
        }
      ]
    },
    "required": [
      "indexFieldName",
      "indexFieldType",
      "dataSourceFieldName"
    ]
  ]
}

```

```

    }
  ]
}
},
"required": [
  "fieldMappings"
]
},
"contacts": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": [
        {
          "type": "object",
          "properties": {
            "indexFieldName": {
              "type": "string"
            },
            "indexFieldType": {
              "type": "string",
              "enum": ["STRING", "STRING_LIST", "DATE"]
            },
            "dataSourceFieldName": {
              "type": "string"
            },
            "dateFieldFormat": {
              "type": "string",
              "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
            }
          }
        },
        "required": [
          "indexFieldName",
          "indexFieldType",
          "dataSourceFieldName"
        ]
      ]
    }
  }
},
"required": [
  "fieldMappings"
]

```

```

    },
    "notes": {
      "type": "object",
      "properties": {
        "fieldMappings": {
          "type": "array",
          "items": [
            {
              "type": "object",
              "properties": {
                "indexFieldName": {
                  "type": "string"
                },
                "indexFieldType": {
                  "type": "string",
                  "enum": ["STRING", "DATE"]
                },
                "dataSourceFieldName": {
                  "type": "string"
                },
                "dateFieldFormat": {
                  "type": "string",
                  "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
                }
              }
            }
          ]
        },
        "required": [
          "indexFieldName",
          "indexFieldType",
          "dataSourceFieldName"
        ]
      }
    },
    "required": [
      "fieldMappings"
    ]
  },
  "required": ["email"]
},
"additionalProperties": {
  "type": "object",

```



```
"properties": {
  "inclusionPatterns": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "exclusionPatterns": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "inclusionUsersList": {
    "type": "array",
    "items": {
      "type": "string",
      "format": "email"
    }
  },
  "exclusionUsersList": {
    "type": "array",
    "items": {
      "type": "string",
      "format": "email"
    }
  },
  "s3bucketName": {
    "type": "string"
  },
  "inclusionUsersFileName": {
    "type": "string"
  },
  "exclusionUsersFileName": {
    "type": "string"
  },
  "inclusionDomainUsers": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "exclusionDomainUsers": {
    "type": "array",
```

```

    "items": {
      "type": "string"
    }
  },
  "crawlCalendar": {
    "type": "boolean"
  },
  "crawlNotes": {
    "type": "boolean"
  },
  "crawlContacts": {
    "type": "boolean"
  },
  "crawlFolderAcl": {
    "type": "boolean"
  },
  "startCalendarDateTime": {
    "anyOf": [
      {
        "type": "string",
        "pattern": "^[0-9]{4}-[0-9]{2}-[0-9]{2}T[0-9]{2}:[0-9]{2}:[0-9]{2}Z$"
      },
      {
        "type": "string",
        "pattern": ""
      }
    ]
  },
  "endCalendarDateTime": {
    "anyOf": [
      {
        "type": "string",
        "pattern": "^[0-9]{4}-[0-9]{2}-[0-9]{2}T[0-9]{2}:[0-9]{2}:[0-9]{2}Z$"
      },
      {
        "type": "string",
        "pattern": ""
      }
    ]
  },
  "subject": {
    "type": "array",
    "items": {
      "type": "string"
    }
  }
}

```

```
    }
  },
  "emailFrom": {
    "type": "array",
    "items": {
      "type": "string",
      "format": "email"
    }
  },
  "emailTo": {
    "type": "array",
    "items": {
      "type": "string",
      "format": "email"
    }
  }
},
"required": [],
"syncMode": {
  "type": "string",
  "enum": [
    "FORCED_FULL_CRAWL",
    "FULL_CRAWL",
    "CHANGE_LOG"
  ]
},
"type" : {
  "type" : "string",
  "pattern": "MSEXCHANGE"
},
"secretArn": {
  "type": "string"
}
},
"version": {
  "type": "string",
  "anyOf": [
    {
      "pattern": "1.0.0"
    }
  ]
},
"required": [
```

```

    "connectionConfiguration",
    "repositoryConfigurations",
    "syncMode",
    "additionalProperties",
    "secretArn",
    "type"
  ]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.
repositoryEndpointMetadata	The endpoint information for the data source.
tenantId	The Microsoft 365 tenant ID. You can find your tenant ID in the Properties of your Azure Active Directory Portal.
repositoryConfigurations	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings.
<ul style="list-style-type: none"> • email • attachment • calendar • contacts • notes 	A list of objects that map the attributes or field names of your Microsoft Exchange data source.
secretARN	The Amazon Resource Name (ARN) of an AWS Secrets Manager secret that contains the key-value pairs required to connect to your Exchange data source. This includes your client ID and your client secret.

Configuration	Description
<code>additionalProperties</code>	Additional configuration options for content in your data source
<code>inclusionPatterns</code> <ul style="list-style-type: none"><code>inclusionUsersList</code><code>inclusionUsersFileName</code><code>inclusionDomainUsers</code>	A list of regular expression patterns to <i>include</i> specific files in your Exchange data source. Files that match the patterns are included in the index. Files that don't match the patterns are excluded from the index. If a file matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence and the file isn't included in the index.
<code>exclusionPatterns</code> <ul style="list-style-type: none"><code>exclusionUsersList</code><code>exclusionUsersFileName</code><code>exclusionDomainUsers</code>	A list of regular expression patterns to <i>exclude</i> specific files in your Exchange data source. Files that match the patterns are excluded from the index. Files that don't match the patterns are included in the index. If a file matches both an exclusion and inclusion pattern, the exclusion pattern takes precedence and the file isn't included in the index.
<code>startCalendarDateTime</code>	Use to specify the date and time for Calendar content to be crawled by Amazon Q.
<code>endCalendarDateTime</code>	Use to specify the date and time for Calendar content to be crawled by Amazon Q.
<code>subject</code>	Use to specify email subject lines to be crawled.
<code>emailFrom</code>	Use to specify emails to be crawled based on sender.
<code>emailTo</code>	Use to specify emails to be crawled based on recipient.

Configuration	Description
<ul style="list-style-type: none"> • <code>crawlCalendar</code> • <code>crawlNotes</code> • <code>crawlContacts</code> • <code>crawlFolderAcl</code> 	<p><code>true</code> to index this content in your Microsoft Exchange data source.</p>
<p><code>syncMode</code></p>	<p>Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose between the following options:</p> <ul style="list-style-type: none"> • Use <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index • Use <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index • Use <code>CHANGE_LOG</code> to incrementally crawl only new and modified content each time your data source syncs with your index
<p><code>type</code></p>	<p>The type of data source. Specify <code>MSEXCHANGE</code> as your data source type.</p>
<p><code>enableIdentityCrawler</code></p>	<p><code>true</code> to activate identity crawler. Identity crawler is activated by default. Crawling identity information on users and groups with access to specific documents is useful for user context filtering. Search results are filtered based on the user or their group access to documents. See Identity crawler for more information.</p>

Configuration	Description
version	The version of this template that's currently supported.

How Amazon Q Business connector crawls Exchange ACLs

When you connect an Exchange data source to Amazon Q Business, Amazon Q Business crawls ACL information attached to a document (user and group information) from your Exchange instance. If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

The Exchange group and user IDs are mapped as follows:

- `_tenant_id` – Your Microsoft tenant ID is a globally unique identifier that's necessary to configure each connector instance. Your tenant ID is different from your organization name or domain and can be found in the properties section of your Microsoft account dashboard.

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business Microsoft Exchange data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q Business enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a attribute mapping already available, or if you want to map additional document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

Important

During Preview, filtering using document attributes in chat is only supported through the API.

Note

You can map any Exchange field to the document title or document body Amazon Q reserved/default index fields.

Supported entities and field mappings

- [Mails](#)
- [Calendar](#)
- [Attachments](#)
- [OneNotes](#)
- [Contacts](#)

Mails

Microsoft Exchange field name	Index field name	Description	Data type
createdDateTime	_created_at	Default	Date

Microsoft Exchange field name	Index field name	Description	Data type
lastModifiedDateTime	_last_updated_at	Default	Date
uri	_source_uri	Default	String
category	_category	Default	String
bccRecipients	xchng_bccRecipient	Custom	String list
ccRecipients	xchng_ccRecipient	Custom	String list
hasAttachment	xchng_hasAttachment	Custom	String
sendDateTime	xchng_sendDateTime	Custom	Date
importance	xchng_importance	Custom	String
from	xchng_from	Custom	String
to	xchng_to	Custom	String list
receivedDateTime	xchng_receivedDateTime	Custom	Date
isRead	xchng_isRead	Custom	String
replyTo	xchng_replyTo	Custom	String
folder	xchng_folder	Custom	String
title	xchng_title	Custom	String
flagStatus	xchng_flagStatus	Custom	String

Calendar

Microsoft Exchange field name	Index field name	Description	Data type
location	xchnng_location	Custom	String
organizer	xchnng_organizer	Custom	String
subject	xchnng_subject	Custom	String
weblink	_source_uri	Default	String
createdDateTime	_created_at	Default	Date
lastModifiedDateTime	_last_updated_at	Default	Date
eventStartTime	xchnng_eventStartTime	Default	Date
eventEndTime	xchnng_eventEndTime	Default	Date
attendees	xchnng_attendees	Custom	String
recurrence	xchnng_Recurrence	Custom	String
category	_category	Default	String
isReminderOn	xchnng_isReminderOn	Custom	String
sensitivity	xchnng_sensitivity	Custom	String
isOnlineMeeting	xchnng_isOnlineMeeting	Custom	String
seriesMasterId	xchnng_seriesMasterId	Custom	String
isCancelled	xchnng_isCancelled	Custom	String

Attachments

Microsoft Exchange field name	Index field name	Description	Data type
title	xchn_g_title	Custom	String
lastModifiedDateTime	_last_updated_at	Default	Date
category	_category	Default	String
contentType	_file_type	Default	String
size	xchn_g_size	Custom	String
url	_source_uri	Default	String

OneNotes

Microsoft Exchange field name	Index field name	Description	Data type
isShared	xchn_g_isShared	Custom	String
link	xchn_g_links	Custom	String
title	xchn_g_title	Custom	String
lastUpdatedBy	xchn_g_lastUpdatedBy	Custom	String
lastModifiedDateTime	_last_updated_at	Default	Date
createdDateTime	_created_at	Default	Date
category	_category	Default	String
createdBy	xchn_g_createdBy	Custom	String

Microsoft Exchange field name	Index field name	Description	Data type
userRole	xchn_g_useRole	Custom	String

Contacts

Microsoft Exchange field name	Index field name	Description	Data type
contactName	xchn_g_contactName	Custom	String
emailAddress	xchn_g_email	Custom	String
companyName	xchn_g_com panyName	Custom	String
manager	xchn_g_manager	Custom	String
jobTitle	xchn_g_jobtitle	Custom	String
location	xchn_g_officeLocation	Custom	String
mobilePhone	xchn_g_mobile	Custom	String
birthday	xchn_g_birthday	Custom	Date
homeAddress	xchn_g_homeAddress	Custom	String
businessAddress	xchn_g_businessAddr ess	Custom	String
department	xchn_g_department	Custom	String
profession	xchn_g_profession	Custom	String
createdAt	_created_at	Default	Date
category	_category	Default	String

Microsoft Exchange field name	Index field name	Description	Data type
url	_source_uri	Custom	String

IAM role for Amazon Q Business Microsoft Exchange connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the `BatchPutDocument` and `BatchDeleteDocument` operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToGetSecret",
      "Effect": "Allow",
      "Action": [
        "secretsmanager:GetSecretValue"
      ],
      "Resource": [
        "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
      ]
    }
  ],
}
```

```

{
  "Sid": "AllowsAmazonQToDecryptSecret",
  "Effect": "Allow",
  "Action": [
    "kms:Decrypt"
  ],
  "Resource": [
    "arn:aws:kms:{{region}}:{{account_id}}:key/[key_id]"
  ],
  "Condition": {
    "StringLike": {
      "kms:ViaService": [
        "secretsmanager.*.amazonaws.com"
      ]
    }
  }
},
{
  "Sid": "AllowsAmazonQToIngestDocuments",
  "Effect": "Allow",
  "Action": [
    "qbusiness:BatchPutDocument",
    "qbusiness:BatchDeleteDocument"
  ],
  "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
},
{
  "Sid": "AllowsAmazonQToIngestPrincipalMapping",
  "Effect": "Allow",
  "Action": [
    "qbusiness:PutGroup",
    "qbusiness:CreateUser",
    "qbusiness>DeleteGroup",
    "qbusiness:UpdateUser",
    "qbusiness:ListGroup"
  ],
  "Resource": [
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}",
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}",
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}/data-source/*"
  ]
}

```

```

},
{
  "Sid": "AllowsAmazonQToCreateAndDeleteNI",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateNetworkInterface",
    "ec2>DeleteNetworkInterface"
  ],
  "Resource": [
    "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
    "arn:aws:ec2:{{region}}:{{account_id}}:security-group/[{{security_group}}]"
  ]
},
{
  "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateNetworkInterface",
    "ec2>DeleteNetworkInterface"
  ],
  "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
  "Condition": {
    "StringLike": {
      "aws:RequestTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
    },
    "ForAllValues:StringEquals": {
      "aws:TagKeys": [
        "AMAZON_Q"
      ]
    }
  }
},
{
  "Sid": "AllowsAmazonQToCreateTags",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateTags"
  ],
  "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
  "Condition": {
    "StringEquals": {
      "ec2:CreateAction": "CreateNetworkInterface"
    }
  }
}

```

```

    },
    {
      "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
      "Effect": "Allow",
      "Action": [
        "ec2:CreateNetworkInterfacePermission"
      ],
      "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
      "Condition": {
        "StringLike": {
          "aws:ResourceTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
        }
      }
    }
  ],
  {
    "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
    "Effect": "Allow",
    "Action": [
      "ec2:DescribeNetworkInterfaces",
      "ec2:DescribeAvailabilityZones",
      "ec2:DescribeNetworkInterfaceAttribute",
      "ec2:DescribeVpcs",
      "ec2:DescribeRegions",
      "ec2:DescribeNetworkInterfacePermissions",
      "ec2:DescribeSubnets"
    ],
    "Resource": "*"
  }
]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
    }
  ]
}

```



```

    "Condition": {
      "StringEquals": {
        "aws:SourceAccount": "{{source_account}}"
      },
      "ArnEquals": {
        "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
      }
    }
  }
]
}

```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Troubleshooting your Amazon Q Business Microsoft Exchange connector

The following table provides information about error codes you may see for the Microsoft Exchange connector and suggested troubleshooting actions.

Error code	Error message	Suggested resolution
MSE-5101	Exception occurred while validating the configurations.	Error occurred while validating the configurations. Verify the configurations and try again.
MSE-5102	Invalid clientId pattern.	Error occurred while validating the configurations. Verify the configurations and try again.
MSE-5103	ClientSecret Over maximum length.	Error occurred while validating the configurations. Verify the configurations and try again.
MSE-5104	Enter valid credentials. Client ID should not be null or empty.	Error occurred while validating the configurations. Client ID should not be null.

Error code	Error message	Suggested resolution
MSE-5105	Enter valid credentials. Client Secret should not be null or empty.	Error occurred while validating the configurations. Client Secret should not be null.
MSE-5106	Enter valid credentials. Tenant ID should not be null or empty	Error occurred while validating the configurations. Tenant ID should not be null.
MSE-5107	The provided client ID is invalid.Please verify the client ID and try again.	Provide client id is invalid while doing authentication. Connection will be unsuccessful. Provide valid client id.
MSE-5108	The provided client secret is invalid. Verify the client secret and try again.	Provide client secret is invalid while doing authentication. Connection will be unsuccessful. Provide valid client secret.
MSE-5109	The provided tenant ID is invalid. Please verify the tenant ID and try again.	Provide tenant ID is invalid while doing authentication. Connection will be unsuccessful. Provide valid tenant ID.
MSE-5200	Got exception from customer while accessing the list of users.	Error occurred while fetching the list of users from Microsoft Graph API. Check logs for more details.
MSE-5201	Got exception from customer while accessing mails.	Error occurred while fetching mails from Microsoft Graph API. Check logs for more details.
MSE-5202	Got exception from customer while accessing calendar events.	Error occurred while fetching calendar events from Microsoft Graph API. Check logs for more details.
MSE-5203	Got exception from customer while accessing OneNotes.	Error occurred while fetching one notes from Microsoft Graph API. Check logs for more details.

Error code	Error message	Suggested resolution
MSE-5204	Got exception from customer while accessing attachments.	Error occurred while fetching attachments from Microsoft Graph API. Check logs for more details.
MSE-5205	Got exception from customer while accessing contacts.	Error occurred while fetching contacts from Microsoft Graph API. Check logs for more details.
MSE-5206	Error occurred while retrying API requests.	Error occurred while retrying API requests to fetch data from Microsoft Graph API.
MSE-5301	Got exception from customer while running changelog mode.	Error occurred while handling changelog token. Refer logs or contact connector team for more information.

Connecting Microsoft OneDrive to Amazon Q Business

Microsoft OneDrive is a cloud-based storage service that you can use to store, share, and host your content. You can connect Microsoft OneDrive instance to Amazon Q Business—using either the AWS Management Console or the [CreateDataSource](#) API—and create an Amazon Q web experience.

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).
- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [Microsoft OneDrive connector overview](#)
- [Prerequisites for connecting Amazon Q Business to Microsoft OneDrive](#)
- [Connecting Amazon Q Business to Microsoft OneDrive using the console](#)

- [Connecting Amazon Q Business to Microsoft OneDrive using APIs](#)
- [How Amazon Q Business connector crawls Microsoft OneDrive ACLs](#)
- [Amazon Q Business Microsoft OneDrive data source connector field mappings](#)
- [IAM role for Amazon Q Business Microsoft OneDrive connector](#)
- [Troubleshooting your Amazon Q Business Microsoft OneDrive connector](#)

Microsoft OneDrive connector overview

The following table gives an overview of the Amazon Q Business Microsoft OneDrive connector and its supported features.

Category	Feature	Support
Security	Authentication type	OAuth 2.0
	Authentication credentials	<ul style="list-style-type: none"> • Microsoft OneDrive Client ID • Microsoft OneDrive Client secret
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Identity crawling	Yes
	VPC	Yes
Crawl features	Custom metadata	No
	Entities	Yes. The following entities are supported: <ul style="list-style-type: none"> • File
	Field mappings	Yes. Supports both default and custom field mappings. For more information, see Field mappings .
	Filters	Yes. The following filters are supported: <ul style="list-style-type: none"> • Include/exclude OneNotes using page name

Category	Feature	Support
		<ul style="list-style-type: none"> • Include/exclude OneNotes using section name • Include/exclude using file path • Include/exclude using file name • Include/exclude using file type
	Sync mode	Supports full and incremental sync.
	File types	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to Microsoft OneDrive

Before you begin, make sure that you have completed the following prerequisites.

In your Azure Active Directory (AD) application, make sure you have:

- Created an Azure Active Directory (AD) application.
- Used the AD application ID to register a secret key for the application on the AD site. The secret key must contain the application ID and a secret key.
- Copied the AD domain of the organization.
- Added the following permissions to your AD application on the Microsoft Graph option:
 - Read files in all site collections (File.Read.All)
 - Read all users' full profiles (User.Read.All)
 - Read all groups (Group.Read.All)
 - Read all notes (Notes.Read.All)

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your Microsoft OneDrive authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Connecting Amazon Q Business to Microsoft OneDrive using the console

The following procedure outlines how to connect Amazon Q Business to Microsoft OneDrive using the AWS Management Console.

Connecting Amazon Q to Microsoft OneDrive

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **Microsoft OneDrive** page, enter the following information:
6. In **Source**, enter the following information:
 - **OneDrive Tenant ID** Enter your OneDrive Tenant ID without the protocol. You can find your OneDrive Tenant ID under Directory ID in the Microsoft Azure AD admin center.
7. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

Note

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

8. In Authentication – Choose between New and Existing.

- If you choose **Existing**, select an existing secret for **Select secret**.

If you choose **New**, enter the following information in the **New AWS Secrets Manager secret** section:

- i. **Secret name** – A name for your secret.
- ii. For **Application ID** and **Application password** – Enter the authentication credential values from your OneDrive account and then choose **Save authentication**.

9. Configure VPC and security group – optional – Choose whether you want to use a VPC. If you do, enter the following information:

- a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
- b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

10. Identity crawler – Choose to activate Amazon Q identity crawler to sync identity information.

For more information, see [Identity crawler](#).

11. IAM role – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

12. In Sync scope, for Select OneDrive users, choose between the following options:

- **Add a user name file** – Choose to add a user names file saved in an Amazon S3 bucket. Provide the path to the file by choosing **Browse**.

Note

If you choose this option, the IAM role for the data source must have read permissions for the Amazon S3 bucket where the file is stored.

- **Add user names here** – You can add a maximum of 10 users using this option. To add more than 10 users, please create a file containing the user names and choose **Add a user name file**.
13. For **Additional configuration – optional**:
- For **Regex for OneNote** and **Regex Patterns** – Add regular expression patterns to include or exclude certain files. You can add up to 100 patterns.
14. For **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.
- **Full sync** – Sync all content regardless of the previous sync status.
 - **New, modified, or deleted content sync** – Only sync new, modified, and deleted content.
15. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
16. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
17. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:
- a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
 - b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

Note

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

18. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

19. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

Note

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to Microsoft OneDrive using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the configuration parameter to provide a JSON schema with all other configuration information specific to your data source connector.

For an example of the API request, see [CreateDataSource](#) in the Amazon Q API Reference.

Microsoft OneDrive JSON schema

The following is a the Microsoft OneDrive JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
      "type": "object",
      "properties": {
        "repositoryEndpointMetadata": {
          "type": "object",
```

```
"properties": {
  "tenantId": {
    "type": "string",
    "pattern": "^[0-9a-f]{8}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{12}$",
    "minLength": 36,
    "maxLength": 36
  }
},
"required": [
  "tenantId"
]
},
"required": [
  "repositoryEndpointMetadata"
],
"repositoryConfigurations": {
  "type": "object",
  "properties": {
    "file": {
      "type": "object",
      "properties": {
        "fieldMappings": {
          "type": "array",
          "items": [
            {
              "type": "object",
              "properties": {
                "indexFieldName": {
                  "type": "string"
                },
                "indexFieldType": {
                  "type": "string",
                  "enum": [
                    "STRING",
                    "STRING_LIST",
                    "DATE",
                    "LONG"
                  ]
                },
                "dataSourceFieldName": {
                  "type": "string"
                }
              }
            }
          ]
        }
      }
    }
  }
}
```

```

        "dateFieldFormat": {
            "type": "string",
            "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
        }
    },
    "required": [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
    ]
}
]
}
},
"required": [
    "fieldMappings"
]
}
}
},
"additionalProperties": {
    "type": "object",
    "properties": {
        "isCrawlAcl": {
            "type": "boolean"
        },
        "fieldForUserId": {
            "type": "string"
        },
        "userNameFilter": {
            "type": "array",
            "items": {
                "type": "string"
            }
        },
        "userFilterPath": {
            "type": "string"
        },
        "isUserNameOnS3": {
            "type": "boolean"
        },
        "inclusionFileTypePatterns": {
            "type": "array",
            "items": {

```

```
    "type": "string"
  }
},
"exclusionFileTypePatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"inclusionFileNamePatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"exclusionFileNamePatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"inclusionFilePathPatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"exclusionFilePathPatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"inclusionOneNoteSectionNamePatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"exclusionOneNoteSectionNamePatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
}
```

```
    },
    "inclusionOneNotePageNamePatterns": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    "exclusionOneNotePageNamePatterns": {
      "type": "array",
      "items": {
        "type": "string"
      }
    }
  },
  "required": []
},

"enableIdentityCrawler": {
  "type": "boolean"
},
"type": {
  "type": "string",
  "pattern": "ONEDRIVEV2"
},
"syncMode": {
  "type": "string",
  "enum": [
    "FULL_CRAWL",
    "FORCED_FULL_CRAWL",
    "CHANGE_LOG"
  ]
},
"secretArn": {
  "type": "string",
  "minLength": 20,
  "maxLength": 2048
}
},
"version": {
  "type": "string",
  "anyOf": [
    {
      "pattern": "1.0.0"
    }
  ]
}
```

```

]
},
"required": [
  "connectionConfiguration",
  "repositoryConfigurations",
  "syncMode",
  "additionalProperties",
  "secretArn",
  "type"
]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.
repositoryEndpointMetadata	The endpoint information for the data source. This includes the tenant ID in the form of the OneDrive site URL.
repositoryConfigurations	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings. Specify the type of data source and the secret ARN.
secretARN	The Amazon Resource Name (ARN) of an AWS Secrets Manager secret that contains the key-value pairs required to connect to your OneDrive. The secret must contain a JSON structure with the following keys: <div data-bbox="829 1667 1507 1866" data-label="Text"> <pre> { "clientId": "OAuth Client ID", "password": "client secret" } </pre> </div>

Configuration	Description
additionalProperties	Additional configuration options for your content in your data source.
isCrawlAcl	Specify true to crawl access control information from documents.
fieldForUserId	Specify field to use for UserId for ACL crawling.
<ul style="list-style-type: none"> • userNameFilter • userFilterPath • inclusionFileTypePatterns • exclusionFileTypePatterns • inclusionFileNamePatterns • exclusionFileNamePatterns • inclusionFilePathPatterns • exclusionFilePathPatterns • inclusionOneNoteSectionNamePatterns • exclusionOneNoteSectionNamePatterns • inclusionOneNotePageNamePatterns • exclusionOneNotepageNamePatterns 	A collection of strings that specifies which entities to filter.

Configuration	Description
isUserNameOnS3	true to provide a list of user names in a file stored in an Amazon S3.
type	The type of data source. Specify ONEDRIVEV2 as your data source type.
enableIdentityCrawler	true to activate identity crawler. Identity crawler is activated by default. Crawling identity information on users and groups with access to specific documents is useful for user context filtering. Search results are filtered based on the user or their group access to documents. See Identity crawler for more information.
syncMode	Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose between the following options: <ul style="list-style-type: none">• Use FORCED_FULL_CRAWL to freshly re-crawl all content and replace existing content each time your data source syncs with your index• Use FULL_CRAWL to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index• Use CHANGE_LOG to incrementally crawl only new and modified content each time your data source syncs with your index

Configuration	Description
version	The version of this template that's currently supported.

How Amazon Q Business connector crawls Microsoft OneDrive ACLs

When you connect an Microsoft OneDrive data source to Amazon Q Business, Amazon Q Business crawls ACL information attached to a document (user and group information) from your Microsoft OneDrive instance. If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

A Microsoft OneDrive data source returns section and page information from OneDrive access control list (ACL) entities. Amazon Q uses the OneDrive tenant domain to connect to the OneDrive instance and can filter based on section name, page type, file name, file type and file contents.

For standard objects, the `_user_id` and `_group_id` are used as follows:

- `_user_id` – Your Microsoft OneDrive user email ID is mapped to the `_user_id` field.
- `_group_id` – Your Microsoft OneDrive group email is mapped to the `_group_id` field.

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business Microsoft OneDrive data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q Business enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.

- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a attribute mapping already available, or if you want to map additional document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q Microsoft OneDrive connector supports the following entities and the associated reserved and custom attributes.

Microsoft OneDrive field name	Index field name	Description	Data type
sourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
lastUpdatedAt	_last_updated_at	Default	Date
createdBy	_authors	Default	String list
lastUpdatedBy	_od_last_updated_by	Custom	String
fs_createdAt	od_fs_created_at	Custom	Date
fs_lastUpdatedAt	of_fs_last_updated_at	Custom	Date

Microsoft OneDrive field name	Index field name	Description	Data type
size	od_size	Custom	Long (numeric)
cTag	od_cTag	Custom	String
eTag	od_eTag	Custom	String
fileMimeType	od_file_mime_type	Custom	String
oneNoteDocument	od_documentName	Custom	String
oneNoteSection	od_sectionName	Custom	String
oneNotePage	od_pageName	Custom	String

IAM role for Amazon Q Business Microsoft OneDrive connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the `BatchPutDocument` and `BatchDeleteDocument` operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
```

```

"Version": "2012-10-17",
"Statement": [
  {
    "Sid": "AllowsAmazonQToGetSecret",
    "Effect": "Allow",
    "Action": [
      "secretsmanager:GetSecretValue"
    ],
    "Resource": [
      "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[[secret_id]]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToDecryptSecret",
    "Effect": "Allow",
    "Action": [
      "kms:Decrypt"
    ],
    "Resource": [
      "arn:aws:kms:{{region}}:{{account_id}}:key/[[key_id]]"
    ],
    "Condition": {
      "StringLike": {
        "kms:ViaService": [
          "secretsmanager.*.amazonaws.com"
        ]
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToIngestDocuments",
    "Effect": "Allow",
    "Action": [
      "qbusiness:BatchPutDocument",
      "qbusiness:BatchDeleteDocument"
    ],
    "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
  },
  {
    "Sid": "AllowsAmazonQToIngestPrincipalMapping",
    "Effect": "Allow",
    "Action": [
      "qbusiness:PutGroup",

```



```

{
  "Sid": "AllowsAmazonQToCreateTags",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateTags"
  ],
  "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
  "Condition": {
    "StringEquals": {
      "ec2:CreateAction": "CreateNetworkInterface"
    }
  }
},
{
  "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateNetworkInterfacePermission"
  ],
  "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
  "Condition": {
    "StringLike": {
      "aws:ResourceTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
    }
  }
},
{
  "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
  "Effect": "Allow",
  "Action": [
    "ec2:DescribeNetworkInterfaces",
    "ec2:DescribeAvailabilityZones",
    "ec2:DescribeNetworkInterfaceAttribute",
    "ec2:DescribeVpcs",
    "ec2:DescribeRegions",
    "ec2:DescribeNetworkInterfacePermissions",
    "ec2:DescribeSubnets"
  ],
  "Resource": "*"
}
]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "{{source_account}}"
        },
        "ArnEquals": {
          "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
        }
      }
    }
  ]
}
```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Troubleshooting your Amazon Q Business Microsoft OneDrive connector

The following table provides information about error codes you may see for the Microsoft OneDrive connector and suggested troubleshooting actions.

Error code	Error message	Suggested resolution
OND-5000	Exception occurred while sending request to OneDrive api for testing connection, please try again later.	Try again.

Error code	Error message	Suggested resolution
OND-5001	Provided client ID key is not Valid.	Provide a valid client ID.
OND-5002	Provided client secret key is not valid.	Provide valid client secret.
OND-5003	Provided tenant ID key is not valid.	Provide a valid tenant ID.
OND-5102	Client ID cannot be null/empty.	Provide a valid client ID.
OND-5103	Tenant ID cannot be null/empty.	Provide a valid tenant Id.
OND-5104	Client Secret cannot be null/empty.	Provide a valid client Secret.
OND-5105	Invalid client ID pattern.	Provide a valid client ID.
OND-5106	Client Secret Over maximum length.	Length of client secret ID should be at least 256. Provide a valid client secret.
OND-5107	User Name Filter/ User Name Path should not be null or empty value.	Provide User Name Filter or User Name Path.
OND-5108	User Name Filter can only support up to 10 users.	Provide up to 10 users in User Name Filter or provide file of list of users in User Name Path.
OND-5109	Users mentioned in the list do not belong to the same domain.	Provide valid list of users which belong to same domain.
OND-5110	Users mentioned in the list are not valid.	Provide valid users.

Error code	Error message	Suggested resolution
OND-5200	Exception occurred while fetching files in full crawl.	Check logs for more details.
OND-5203	Exception occurred while fetching drive files.	Provide correct credentials.
OND-5204	Exception occurred while fetching OneNote files.	Check logs for more details.
OND-5300	Exception occurred while fetching files in change log.	Check logs for more details.
OND-5400	Exception occurred while building group details.	Check logs for more details.
OND-5401	Exception occurred while fetching list of groups.	Check logs for more details.
OND-5500	Exception occurred while getting file content response.	Check logs for more details.
OND-5501	Only String, String List, Date and Long formats are supported for field mappings.	Please provide valid formats in field mappings.
OND-5502	Exception occurred while fetching OneNote files.	Check logs for more details.

Connecting SharePoint (Online) to Amazon Q Business

Microsoft SharePoint is a collaborative website building service that lets you customize web content and create web pages, web sites, document libraries, and lists. You can connect SharePoint

(Online) instance to Amazon Q Business—using either the AWS Management Console or the [CreateDataSource](#) API—and create an Amazon Q web experience.

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).
- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [SharePoint \(Online\) connector overview](#)
- [Prerequisites for connecting Amazon Q Business to SharePoint \(Online\)](#)
- [Connecting Amazon Q Business to SharePoint \(Online\) using the console](#)
- [Connecting Amazon Q Business to SharePoint \(Online\) using APIs](#)
- [How Amazon Q Business connector crawls SharePoint \(Online\) ACLs](#)
- [Amazon Q Business SharePoint \(Online\) data source connector field mappings](#)
- [IAM role for Amazon Q SharePoint \(Online\) connector](#)
- [Known limitations for the Amazon Q Business SharePoint \(Online\) connector](#)
- [Troubleshooting your Amazon Q Business SharePoint \(Online\) connector](#)

SharePoint (Online) connector overview

The following table gives an overview of the Amazon Q Business SharePoint (Online) connector and its supported features.

Category	Feature	Support
Security	Authentication type	Basic, OAuth 2.0, Azure AD App-Only, SharePoint App-Only
	Authentication credentials	Basic <ul style="list-style-type: none"> • SharePoint (Online) admin username

Category	Feature	Support
		<ul style="list-style-type: none">• SharePoint (Online) admin password <p>OAuth 2.0</p> <ul style="list-style-type: none">• SharePoint Tenant ID• SharePoint admin username• SharePoint admin password• Client ID• Client secret <p>Azure App-Only</p> <ul style="list-style-type: none">• Tenant ID• Certificate path• Client ID• Private key <p>SharePoint App-Only</p> <ul style="list-style-type: none">• Tenant ID• Azure AD Client ID• Azure AD Client secret• SharePoint App-Only Client ID• SharePoint App-Only Client secret <p>OAuth 2.0 refresh token</p> <ul style="list-style-type: none">• Tenant ID• Azure AD Client ID• Azure AD Client secret• Refresh token

Category	Feature	Support
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Integration with Identity Provider (IdP)	Yes. Azure AD.
	Identity crawling	Yes
	VPC	Yes
Crawl features	Custom metadata	Yes. Supports custom metadata for File entity only.
	Entities	Yes. The following entities are supported: <ul style="list-style-type: none"> • Files • Attachments • Link • Pages • Events • Comments
	Field mappings	Yes. Supports both default and custom field mappings. For more information, see Field mappings .

Category	Feature	Support
	Filters	<p>Yes. The following filters are supported:</p> <ul style="list-style-type: none"> • Include/exclude by Links • Include/exclude by Pages • Include/exclude by Events • Include/exclude by file name • Include/exclude by file path • Include/exclude by file type • Include/exclude by OneNote Section name • Include/exclude by OneNote Page name
	<u>Sync mode</u>	Supports full and incremental sync.
	<u>File types</u>	Supports all files supported by Amazon Q.
	<u>Crawled as a document</u>	<ul style="list-style-type: none"> • Each event • Each page • Each file • Each link • Each file attachment • Each comment • Each OneNote

Prerequisites for connecting Amazon Q Business to SharePoint (Online)

The following page outlines the prerequisites you need to complete before connecting SharePoint (Online) to Amazon Q, based on the authentication mode of your choice.

Topics

- [Prerequisites for using basic authentication](#)
- [Prerequisites for using OAuth 2.0 authentication](#)
- [Prerequisites for using Azure AD App-Only authentication](#)

- [Prerequisites for using SharePoint App-Only authentication](#)
- [Prerequisites for using OAuth 2.0 refresh token authentication](#)

Prerequisites for using basic authentication

If you're using basic authentication, make sure you've completed the following steps in SharePoint (Online):

- Copied your SharePoint (Online) instance URLs. The format for the host URL you enter is *https://yourdomain.sharepoint.com/sites/mysite*. Your URL must start with https and contain sharepoint.com.
- Copied the domain name of your SharePoint (Online) instance URL.
- Noted your basic authentication credentials containing the username and password that you use to connect to SharePoint (Online) Online.
- Deactivated **Security Defaults** in your Azure portal using an administrative user. For more information on managing security default settings in the Azure portal, see [Microsoft documentation on how to enable/disable security defaults](#).
- Deactivated multi-factor authentication (MFA) in your SharePoint account, so that Amazon Q is not blocked from crawling your SharePoint content.

Note

No API permissions are required for crawling entities using **Basic authentication**.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your SharePoint (Online) authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Prerequisites for using OAuth 2.0 authentication

If you're using OAuth 2.0 authentication, make sure you've completed the following steps in SharePoint (Online):

- Copied your SharePoint (Online) instance URLs. The format for the host URL you enter is *https://yourdomain.sharepoint.com/sites/mysite*. Your URL must start with `https` and contain `sharepoint.com`.
- Copied the domain name of your SharePoint (Online) instance URL.
- Copied the tenant ID of your Microsoft SharePoint (Online) instance. For details on how to find your tenant ID, see [Find your Microsoft 365 tenant ID](#) on the Microsoft website.
- Noted the username and password that you use to connect to SharePoint (Online).
- Noted the Client ID and Client secret generated after registering SharePoint (Online) with Azure AD.
- **If you're *not* using ACL**, added the following permissions:

Microsoft Graph	SharePoint
<ul style="list-style-type: none"> • Notes.Read.All (Application) – Read all OneNote notebooks • Sites.Read.All (Application) – Read items in all site collections 	<ul style="list-style-type: none"> • AllSites.Read (Delegated) – Read items in all site collections

Note

Note.Read.All and Sites.Read.All are required only if you want to crawl OneNote Documents.

- **If you're using ACL**, added the following permissions:

Microsoft Graph

- GroupMember.Read.All (Application) – Read all group memberships
- Notes.Read.All (Application) – Read all OneNote notebooks
- Sites.FullControl.All (Delegated) – Have full control of all site collections
- Sites.Read.All (Application) – Read items in all site collections
- User.Read.All (Application) – Read all users' full profiles

SharePoint

- AllSites.Read (Delegated) – Read items in all site collections

Note

GroupMember.Read.All and User.Read.All are required only if **Identity crawler** is activated.

- Deactivated multi-factor authentication (MFA) in your SharePoint account, so that Amazon Q is not blocked from crawling your SharePoint content.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your SharePoint (Online) authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Prerequisites for using Azure AD App-Only authentication

If you're using Azure AD App-Only authentication, make sure you've completed the following steps in SharePoint (Online):

- Copied your SharePoint (Online) instance URLs. The format for the host URL you enter is *https://yourdomain.sharepoint.com/sites/mysite*. Your URL must start with `https` and contain `sharepoint.com`.
- Copied the domain name of your SharePoint (Online) instance URL.
- Copied the tenant ID of your Microsoft SharePoint (Online) instance. For details on how to find your tenant ID, see [Find your Microsoft 365 tenant ID](#) on the Microsoft website.
- Noted the file path to a X.509 certificate you have created and stored in an Amazon S3 bucket.
- Noted the private key and the Client ID you generated after registering SharePoint (Online) with Azure AD.
- **If you're *not* using ACL, added the following permissions:**

SharePoint

- Sites.Read.All (Application) – Read items in all site collections
- **If you're using ACL, added the following permissions:**

SharePoint

- Sites.FullControl.All (Application) – Have full control of all site collections

Note

If you want to crawl specific sites, you can choose to restrict permissions to specific sites rather than all sites available in the domain. To do this, use the Sites.Selected (Application) permission. With this API permission, you need to set access permission

on every site explicitly through the Microsoft Graph API. For more information, see [Microsoft's blog on Sites.Selected permissions](#).

- Deactivated multi-factor authentication (MFA) in your SharePoint account, so that Amazon Q is not blocked from crawling your SharePoint content.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your SharePoint (Online) authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Prerequisites for using SharePoint App-Only authentication

If you're using SharePoint App-Only authentication, make sure you've completed the following steps in SharePoint (Online):

- Copied your SharePoint (Online) instance URLs. The format for the host URL you enter is *https://yourdomain.sharepoint.com/sites/mysite*. Your URL must start with https and contain sharepoint.com.
- Copied the domain name of your SharePoint (Online) instance URL.
- Copied the tenant ID of your Microsoft SharePoint (Online) instance. For details on how to find your tenant ID, see [Find your Microsoft 365 tenant ID](#) on the Microsoft website.
- Noted your SharePoint (Online) client ID and client secret generated while granting permission to SharePoint App-Only, and your Client ID and Client secret generated when you registered your SharePoint (Online) app with Azure AD.

- **If you're crawling OneNote documents and using Identity crawler**, added the following permissions:

Microsoft Graph

- GroupMember.Read.All (Application) – Read all group memberships
- Notes.Read.All (Application) – Read all OneNote notebooks
- Sites.Read.All (Application) – Read items in all site collections
- User.Read.All (Application) – Read all users' full profiles

Note

No API permissions are required for crawling entities using SharePoint (Online) **App-Only authentication**.

- Deactivated multi-factor authentication (MFA) in your SharePoint account, so that Amazon Q is not blocked from crawling your SharePoint content.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your SharePoint (Online) authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Prerequisites for using OAuth 2.0 refresh token authentication

If you're using OAuth 2.0 refresh token authentication, make sure you've completed the following steps in SharePoint (Online):

- Copied your SharePoint (Online) instance URLs. The format for the host URL you enter is *https://yourdomain.sharepoint.com/sites/mysite*. Your URL must start with `https` and contain `sharepoint.com`.
- Copied the domain name of your SharePoint (Online) instance URL.
- Copied the tenant ID of your Microsoft SharePoint (Online) instance. For details on how to find your tenant ID, see [Find your Microsoft 365 tenant ID](#) on the Microsoft website.
- Noted your SharePoint (Online) client ID and client secret generated while granting permission to SharePoint App-Only, and your Client ID and Client secret generated when you registered your SharePoint (Online) app with Azure AD.
- **If you're *not* using ACL**, added the following permissions:

Microsoft Graph	SharePoint
<ul style="list-style-type: none"> • Notes.Read.All (Application) – Read all OneNote notebooks • Sites.Read.All (Application) – Read items in all site collections 	<ul style="list-style-type: none"> • AllSites.Read (Delegated) – Read items in all site collections

Note

Note.Read.All and Sites.Read.All are required only if you want to crawl OneNote Documents.

- **If you're using ACL**, added the following permissions:

Microsoft Graph

- GroupMember.Read.All (Application) – Read all group memberships
- Notes.Read.All (Application) – Read all OneNote notebooks
- Sites.FullControl.All (Delegated) – Have full control of all site collections
- Sites.Read.All (Application) – Read items in all site collections
- User.Read.All (Application) – Read all users' full profiles

SharePoint

- AllSites.Read (Delegated) – Read items in all site collections

Note

GroupMember.Read.All and User.Read.All are required only if **Identity crawler** is activated.

- **If you're crawling OneNote documents and using Identity crawler**, added the following permissions:

Microsoft Graph

- GroupMember.Read.All (Application) – Read all group memberships
- Notes.Read.All (Application) – Read all OneNote notebooks
- Sites.Read.All (Application) – Read items in all site collections
- User.Read.All (Application) – Read all users' full profiles

Note

No API permissions are required for crawling entities using SharePoint (Online) **App-Only authentication**.

- Deactivated multi-factor authentication (MFA) in your SharePoint account, so that Amazon Q is not blocked from crawling your SharePoint content.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your SharePoint (Online) authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Connecting Amazon Q Business to SharePoint (Online) using the console

The following procedure outlines how to connect Amazon Q Business to SharePoint (Online) using the AWS Management Console.

Connecting Amazon Q to SharePoint (Online)

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).

4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **SharePoint (Online)** page, enter the following information:

6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. In **Source**, enter the following information:

- a. In **Source**, for **Hosting Method** – Choose **SharePoint Online**.

- b. **Site URLs specific to your SharePoint repository** – Enter the SharePoint host URLs. The format for the host URLs you enter is *https://yourcompany.sharepoint.com/sites/mysite*. The URL must start with https protocol. Separate URLs with a new line. You can add up to 100 URLs.

- c. **Domain** – Enter the SharePoint domain. For example, the domain in the URL *https://yourdomain.sharepoint.com/sites/mysite* is *yourdomain*.

8. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

Note

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

Important

If you don't specify a value, **Email** is considered as the default value.

9. For **Authentication**, choose between **Basic**, **OAuth 2.0**, **Azure AD App-Only authentication**, **SharePoint App-Only authentication**, and **OAuth 2.0 refresh token authentication** based on your use case.

Note

OneNote can only be crawled by the connector using a Tenant ID, and with OAuth 2.0, OAuth 2.0 refresh token, or SharePoint (Online) App Only authentication activated.

- a. If using **Basic Authentication**, enter the following information:
 - For **AWS Secrets Manager secret** – Choose an existing secret or create a Secrets Manager secret to store your SharePoint authentication credentials. If you choose to create a secret, an AWS Secrets Manager secret window opens. Enter the following information in the window:
 - **Secret name** – A name for your secret.
 - **Username** – Username for your SharePoint account.
 - **Password** – Password for your SharePoint account.
- b. If using **OAuth 2.0 authentication**, enter the following information:
 - **Tenant ID** – Tenant ID of your SharePoint account.
 - For **AWS Secrets Manager secret** – Choose an existing secret or create a Secrets Manager secret to store your SharePoint authentication credentials. If you choose to create a secret, an AWS Secrets Manager secret window opens. Enter the following information in the window:
 - **Secret name** – A name for your secret.
 - **Username** – Username for your SharePoint account.
 - **Password** – Password for your SharePoint account.
 - **Client ID** – The Azure AD client ID generated when you register SharePoint in Azure AD.
 - **Client secret** – The Azure AD client secret generated when you register SharePoint in Azure AD.
- c. If using **Azure AD App-Only authentication**, enter the following information:
 - **Tenant ID** – Tenant ID of your SharePoint account.
 - **Azure AD self-signed X.509 certificate** – Certificate to authenticate the connector for Azure AD.

- For **AWS Secrets Manager secret** – Choose an existing secret or create a Secrets Manager secret to store your SharePoint authentication credentials. If you choose to create a secret, an AWS Secrets Manager secret window opens. Enter the following information in the window:
 - **Secret name** – A name for your secret.
 - **Client ID** – The Azure AD client ID generated when you register SharePoint in Azure AD.
 - **Private key** – A private key to authenticate the connector for Azure AD.
- d. If using **SharePoint App-Only authentication**, enter the following information:
 - **Tenant ID**–Tenant ID of your SharePoint account.
 - For **AWS Secrets Manager secret** — Choose an existing secret or create a Secrets Manager secret to store your SharePoint authentication credentials. If you choose to create a secret, an AWS Secrets Manager secret window opens. Enter the following information in the window:
 - **Secret name** – A name for your secret.
 - **SharePoint client ID** – The SharePoint client ID you generated when you registered App-Only at Tenant Level. ClientID format is *ClientID@TenantId*. For example, *ffa956f3-8f89-44e7-b0e4-49670756342c@888d0b57-69f1-4fb8-957f-e1f0bedf82fe*.
 - **SharePoint client secret** – The SharePoint client secret generated when you register for App-Only at Tenant Level.
 - **Client ID** – The Azure AD client ID generated when you register SharePoint in Azure AD.
 - **Client secret** – The Azure AD client secret generated when you register SharePoint to Azure AD.
- e. If using **OAuth 2.0 refresh token authentication**, enter the following information:
 - **Tenant ID**–Tenant ID of your SharePoint account.
 - For **AWS Secrets Manager secret** — Choose an existing secret or create a Secrets Manager secret to store your SharePoint authentication credentials. If you choose to create a secret, an AWS Secrets Manager secret window opens. Enter the following information in the window:
 - **Secret name** – A name for your secret.

- **Client ID** – The Azure AD client ID generated when you register SharePoint in Azure AD.
 - **Client secret** – The Azure AD client secret generated when you register SharePoint to Azure AD.
 - **Refresh token** – The refresh token you've generated to connect SharePoint (Online) to Amazon Q.
10. **Configure VPC and security group – optional** – Choose whether you want to use a VPC. If you do, enter the following information:
- a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
 - b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

11. **Identity crawler** – Choose to activate Amazon Q identity crawler to sync identity information. Only **Local Group Members** will be crawled using **Identity crawler**. For more information, see [Identity crawler](#).

 **Note**

Crawl AD Group mapping is available only for OAuth 2.0, OAuth 2.0 refresh token, and SharePoint (Online) App-Only authentication.

12. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

13. In **Sync scope**, choose from the following options :
- a. **Select entities** – Choose the entities that you want to crawl. You can select to crawl **All** entities or any combination of **Files, Attachments, Links, Pages, Events, Comments, and List Data**.

- b. In **Additional configuration – optional**, for **Entity regex patterns** – Add regular expression patterns for **Links**, **Pages**, and **Events** to include specific entities instead of syncing all your documents.
- c. In **Additional configuration**, for **Regex patterns** – Add regular expression patterns to include or exclude files by **File path**, **File name**, **File type**, **OneNote section name**, and **OneNote page name** instead of syncing all your documents. You can add up to 100 patterns.

 **Note**

OneNote crawling is available only for OAuth 2.0, OAuth 2.0 refresh token, and SharePoint App Only authentication.

14. In **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.
 - **Full sync** – Sync all content regardless of the previous sync status.
 - **New or modified content sync** – Sync only new and modified documents.
 - **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.

For more details, see [Sync mode](#).

15. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
16. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
17. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:
 - a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
 - b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

Note

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

18. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

19. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

Note

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to SharePoint (Online) using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the `configuration` parameter to provide a JSON schema with all other configuration information specific to your data source connector.

For an example of the API request, see [CreateDataSource](#) in the Amazon Q API Reference.

Microsoft SharePoint JSON schema

The following is the Microsoft SharePoint JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
```

```
"properties": {
  "connectionConfiguration": {
    "type": "object",
    "properties": {
      "repositoryEndpointMetadata": {
        "type": "object",
        "properties": {
          "tenantId": {
            "type": "string",
            "pattern": "^[0-9a-f]{8}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{12}$",
            "minLength": 36,
            "maxLength": 36
          },
          "domain": {
            "type": "string"
          }
        },
        "siteUrls": {
          "type": "array",
          "items": {
            "type": "string",
            "pattern": "https://.*"
          }
        }
      },
      "repositoryAdditionalProperties": {
        "type": "object",
        "properties": {
          "s3bucketName": {
            "type": "string"
          },
          "s3certificateName": {
            "type": "string"
          },
          "authType": {
            "type": "string",
            "enum": [
              "OAuth2",
              "OAuth2Certificate",
              "OAuth2App",
              "OAuth2_RefreshToken",
              "Basic",
              "NTLM",
              "Kerberos"
            ]
          }
        }
      }
    }
  }
}
```

```

    "version": {
      "type": "string",
      "enum": [
        "Server",
        "Online"
      ]
    },
    "onPremVersion": {
      "type": "string",
      "enum": [
        "",
        "2013",
        "2016",
        "2019",
        "SubscriptionEdition"
      ]
    }
  },
  "required": [
    "authType",
    "version"
  ]
},
"required": [
  "siteUrls",
  "domain",
  "repositoryAdditionalProperties"
]
},
"required": [
  "repositoryEndpointMetadata"
]
},
"repositoryConfigurations": {
  "type": "object",
  "properties": {
    "event": {
      "type": "object",
      "properties": {
        "fieldMappings": {
          "type": "array",
          "items": [

```

```
{
  "type": "object",
  "properties": {
    "indexFieldName": {
      "type": "string"
    },
    "indexFieldType": {
      "type": "string",
      "enum": [
        "STRING",
        "STRING_LIST",
        "DATE"
      ]
    },
    "dataSourceFieldName": {
      "type": "string"
    },
    "dateFieldFormat": {
      "type": "string",
      "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
    }
  },
  "required": [
    "indexFieldName",
    "indexFieldType",
    "dataSourceFieldName"
  ]
}

]
}
},
"required": [
  "fieldMappings"
]
},
"page": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": [
        {
          "type": "object",
          "properties": {
```

```
    "indexFieldName": {
      "type": "string"
    },
    "indexFieldType": {
      "type": "string",
      "enum": [
        "STRING",
        "DATE",
        "LONG"
      ]
    },
    "dataSourceFieldName": {
      "type": "string"
    },
    "dateFieldFormat": {
      "type": "string",
      "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
    }
  },
  "required": [
    "indexFieldName",
    "indexFieldType",
    "dataSourceFieldName"
  ]
}
]
}
},
"required": [
  "fieldMappings"
]
},
"file": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": [
        {
          "type": "object",
          "properties": {
            "indexFieldName": {
              "type": "string"
            }
          }
        }
      ]
    }
  }
}
```



```

    "indexFieldType": {
      "type": "string",
      "enum": [
        "STRING",
        "DATE",
        "LONG"
      ]
    },
    "dataSourceFieldName": {
      "type": "string"
    },
    "dateFieldFormat": {
      "type": "string",
      "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
    }
  ],
  "required": [
    "indexFieldName",
    "indexFieldType",
    "dataSourceFieldName"
  ]
}
]
}
},
"required": [
  "fieldMappings"
]
},
"link": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": [
        {
          "type": "object",
          "properties": {
            "indexFieldName": {
              "type": "string"
            },
            "indexFieldType": {
              "type": "string",
              "enum": [

```



```

    ]
  },
  "dataSourceFieldName": {
    "type": "string"
  },
  "dateFieldFormat": {
    "type": "string",
    "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
  }
},
"required": [
  "indexFieldName",
  "indexFieldType",
  "dataSourceFieldName"
]
}
]
}
},
"required": [
  "fieldMappings"
]
},
"comment": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": [
        {
          "type": "object",
          "properties": {
            "indexFieldName": {
              "type": "string"
            },
            "indexFieldType": {
              "type": "string",
              "enum": [
                "STRING",
                "STRING_LIST",
                "DATE"
              ]
            }
          }
        }
      ]
    }
  }
},
"dataSourceFieldName": {

```

```

        "type": "string"
      },
      "dateFieldFormat": {
        "type": "string",
        "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
      }
    },
    "required": [
      "indexFieldName",
      "indexFieldType",
      "dataSourceFieldName"
    ]
  }
]
}
},
"required": [
  "fieldMappings"
]
}
}
},
"additionalProperties": {
  "type": "object",
  "properties": {
    "eventTitleFilterRegEx": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    "pageTitleFilterRegEx": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    "linkTitleFilterRegEx": {
      "type": "array",
      "items": {
        "type": "string"
      }
    }
  }
},
"inclusionFilePath": {

```

```
"type": "array",
  "items": {
    "type": "string"
  }
},
"exclusionFilePath": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"inclusionFileTypePatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"exclusionFileTypePatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"inclusionFileNamePatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"exclusionFileNamePatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"inclusionOneNoteSectionNamePatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"exclusionOneNoteSectionNamePatterns": {
  "type": "array",
  "items": {
```

```
    "type": "string"
  }
},
"inclusionOneNotePageNamePatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"exclusionOneNotePageNamePatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"crawlFiles": {
  "type": "boolean"
},
"crawlPages": {
  "type": "boolean"
},
"crawlEvents": {
  "type": "boolean"
},
"crawlComments": {
  "type": "boolean"
},
"crawlLinks": {
  "type": "boolean"
},
"crawlAttachments": {
  "type": "boolean"
},
"crawlListData": {
  "type": "boolean"
},
"crawlAcl": {
  "type": "boolean"
},
"aclConfiguration": {
  "type": "string",
  "enum": [
    "ACLWithLDAPEmailFmt",
    "ACLWithManualEmailFmt",
```

```
    "ACLWithUsernameFmt"
  ]
},
"emailDomain": {
  "type": "string"
},
"isCrawlLocalGroupMapping": {
  "type": "boolean"
},
"isCrawlAdGroupMapping": {
  "type": "boolean"
},
"proxyHost": {
  "type": "string"
},
"proxyPort": {
  "type": "string"
}
},
"required": [
]
},
"type": {
  "type": "string",
  "pattern": "SHAREPOINTV2"
},
"enableIdentityCrawler": {
  "type": "boolean"
},
"syncMode": {
  "type": "string",
  "enum": [
    "FULL_CRAWL",
    "FORCED_FULL_CRAWL",
    "CHANGE_LOG"
  ]
},
"secretArn": {
  "type": "string",
  "minLength": 20,
  "maxLength": 2048
}
},
"version": {
```

```

"type": "string",
"anyOf": [
  {
    "pattern": "1.0.0"
  }
]
},
"required": [
  "connectionConfiguration",
  "repositoryConfigurations",
  "enableIdentityCrawler",
  "syncMode",
  "additionalProperties",
  "secretArn",
  "type"
]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.
repositoryEndpointMetadata	The endpoint information for the data source.
tenantId	The tenant id of your SharePoint account.
domain	The domain of your SharePoint account.
siteUrls	The host URLs of your SharePoint account.
repositoryAdditionalProperties	Additional properties to connect with your repository endpoint.
s3bucketName	The name of the Amazon S3 bucket that stores your Azure AD self-signed X.509 certificate.

Configuration	Description
s3certificateName	The name of the SSL certificate stored in your Amazon S3 bucket.
authType	The type of authentication you are using: OAuth2, OAuth2Certificate , OAuth2App , OAuth2_RefreshToken or Basic.
version	The SharePoint version you are using: Online.
onPremVersion	Not required if you are using SharePoint Online.
repositoryConfigurations	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings.
<ul style="list-style-type: none"> • event • page • file • link • attachment • comment 	A list of objects that map the attributes or field names of your SharePoint (Online) pages and assets to Amazon Q index field names.
additionalProperties	Additional configuration options for your content in your data source.

Configuration	Description
<ul style="list-style-type: none">• eventTitleFilterRegEx• pageTitleFilterRegEx• linkTitleFilterRegEx• inclusionFilePath• exclusionFilePath• inclusionFileTypePatterns• exclusionFileTypePatterns• inclusionFileNamePatterns• exclusionFileNamePatterns• inclusionOneNoteSectionNamePatterns• exclusionOneNoteSectionNamePatterns• inclusionOneNotePageNamePatterns• exclusionOneNotePageNamePatterns• aclConfiguration• emailDomain• proxyHost• proxyPort	<p>A list of regular expression patterns to include/exclude specific files in your SharePoint data source. Files that match the patterns are included in the index. File that don't match the patterns are excluded from the index. If a file matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence, and the file isn't included in the index.</p>

Configuration	Description
<ul style="list-style-type: none"> • <code>crawlFiles</code> • <code>crawlPages</code> • <code>crawlEvents</code> • <code>crawlComments</code> • <code>crawlLinks</code> • <code>crawlAttachments</code> • <code>crawlListData</code> • <code>crawlAcl</code> • <code>isCrawlLocalGroupMapping</code> • <code>isCrawlAdGroupMapping</code> 	<p>Input TRUE to index.</p>
<p><code>type</code></p>	<p>Specify <code>SHAREPOINTV2</code> as your data source type</p>
<p><code>enableIdentityCrawler</code></p>	<p>true to activate identity crawler. Identity crawler is activated by default. Crawling identity information on users and groups with access to specific documents is useful for user context filtering. Search results are filtered based on the user or their group access to documents. See Identity crawler for more information.</p>

Configuration	Description
syncMode	<p>Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose between the following options:</p> <ul style="list-style-type: none"> • Use <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index • Use <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index • Use <code>CHANGE_LOG</code> to incrementally crawl only new and modified content each time your data source syncs with your index
secretARN	<p>The Amazon Resource Name (ARN) of an AWS Secrets Manager secret that contains the key-value pairs required to connect to your SharePoint. If you use basic authentication provide the username and password. If you use OAuth 2.0 authentication, provide the username, password, client ID, and client secret.</p>
version	<p>The version of this template that's currently supported.</p>

How Amazon Q Business connector crawls SharePoint (Online) ACLs

When you connect an SharePoint (Online) data source to Amazon Q Business, Amazon Q crawls ACL information attached to a document (user and group information) from your SharePoint (Online) instance. If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

To filter using a username, use the **User principal name** from your Azure portal. For example, johnstiles@kendra.onmicrosoft.com.

When you use a SharePoint group for user context filtering, calculate the group ID as follows:

For local groups

1. Get the site name. For example, `https://host.onmicrosoft.com/sites/siteName`.
2. Take the SHA256 hash of the site name. For example, `430a6b90503eef95c89295c8999c7981`.
3. Create the group ID by concatenating the SHA256 hash with a vertical bar (|) and the group name. For example, if the group name is "local group name", the group ID is the following:

`"430a6b90503eef95c89295c8999c7981 | localGroupName"` (with a space before and after the vertical bar).

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business SharePoint (Online) data source connector field mappings

To help you structure data for retrieval and chat filtering, Amazon Q Business crawls data source document attributes or metadata and maps them to fields in your Amazon Q index.

Amazon Q has reserved fields that it uses when querying your application. When possible, Amazon Q automatically maps these built-in fields to attributes in your data source. If a built-in field doesn't have a default mapping, or if you want to map additional index fields, use the custom field mappings to specify how a data source attribute maps to your Amazon Q application. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

⚠ Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q Sharepoint connector supports the following entities and the associated reserved and custom attributes.

ℹ Note

You can map any Sharepoint field to the document title or document body Amazon Q reserved/default index fields.

Supported entities and field mappings

- [Files](#)
- [Events](#)
- [Pages](#)
- [Links](#)
- [Attachments](#)
- [Comments](#)

Files

Sharepoint field name	Index field name	Description	Data type
title	sp_title	Custom	String
sourceUri	_source_uri	Default	String
checkInComment	sp_checkInComment	Custom	String
size	sp_sizeLong	Custom	Long (numeric)

Sharepoint field name	Index field name	Description	Data type
lastModifiedDateTime	_last_updated_at	Default	Date
createdAt	_created_at	Default	Date
author	_authors	Default	String list
majorVersion	sp_majorVersion	Custom	String
uiVersionLabel	sp_uiVersionLabel	Custom	String
uniqueId	sp_uniqueId	Custom	String
irmEnabled	sp_irmEnabled	Custom	String
checkOutType	sp_checkOutType	Custom	String
category	_category	Default	String
modifiedBy	sp_modifiedBy	Custom	String
level	sp_level	Custom	String
uiVersion	sp_uiVersion	Custom	String
contentTag	sp_contentTag	Custom	String
eTag	sp_eTag	Custom	String
oneNoteDocument	sp_oneNoteDocument	Custom	String
oneNoteSection	sp_oneNoteSection	Custom	String
oneNotePage	sp_oneNotePage	Custom	String

Events

Sharepoint field name	Index field name	Description	Data type
title	sp_title	Custom	String
lastModifiedDateTime	_last_updated_at	Default	Date
sourceUri	_source_uri	Default	String
attachments	sp_hasAttachments	Custom	String
createdDate	_created_at	Default	Date
authorId	sp_authorId	Custom	String
editorId	sp_editorId	Custom	String
location	sp_location	Custom	String
eventDate	sp_eventDate	Custom	Date
eventEndDate	sp_eventEndDate	Custom	Date
ifRecurrence	sp_ifRecurrence	Custom	String
ifAllDayEvent	sp_ifAllDayEvent	Custom	String
category	_category	Default	String
eventCategory	sp_eventcategory	Custom	String

Pages

Sharepoint field name	Index field name	Description	Data type
createdDateTime	_created_at	Default	Date

Sharepoint field name	Index field name	Description	Data type
lastModifiedDateTime	_last_updated_at	Default	Date
title	sp_title	Custom	String
sourceUri	_source_uri	Default	String
firstPublishedDate	sp_firstPublishedDate	Custom	Date
authorId	sp_authorId	Custom	String
editorId	sp_editorId	Custom	String
category	_category	Default	String

Links

Sharepoint field name	Index field name	Description	Data type
createdAt	_created_at	Default	Date
lastModifiedDateTime	_last_updated_at	Default	Date
title	sp_title	Custom	String
sourceUri	_source_uri	Default	String
fileType	sp_fileType	Custom	String
fileDirPath	sp_fileDirPath	Custom	String
firstPublishedDate	sp_firstPublishedDate	Custom	Date
authorId	sp_authorId	Custom	String

Sharepoint field name	Index field name	Description	Data type
editorId	sp_editorId	Custom	String
category	_category	Default	String
size	sp_sizeLong	Custom	Long (numeric)

Attachments

Sharepoint field name	Index field name	Description	Data type
title	sp__title	Custom	String
parentCreatedDate	_created_at	Default	Date
sourceUri	_source_uri	Default	String
parentModifiedDate	_last_updated_at	Custom	Date
parentListId	sp_parentListId	Custom	String
parentTitle	sp_parentTitle	Custom	String
category	_category	Default	String

Comments

Sharepoint field name	Index field name	Description	Data type
createdDateTime	_created_at	Default	Date
likedBy	sp_likedBy	Custom	String
sourceUri	_source_uri	Custom	String

Sharepoint field name	Index field name	Description	Data type
isReply	sp_isReply	Custom	String
author	_authors	Default	String list
listId	sp_listId	Custom	String
category	_category	Default	String
replyCount	sp_replyCount	Custom	String
parentTitle	sp_parentTitle	Custom	String

IAM role for Amazon Q SharePoint (Online) connector

Note

(Optional) If you use **Azure App-Only authentication**, you also need to add permissions for Amazon Q to access the certificate stored in your Amazon S3 bucket.

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the BatchPutDocument and BatchDeleteDocument operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.

- Permission to access the SSL certificate stored in your Amazon S3 bucket.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AllowsAmazonQToGetS3Objects",
    "Action": [
      "s3:GetObject"
    ],
    "Resource": [
      "arn:aws:s3:::{{input_bucket_name}}/*"
    ],
    "Effect": "Allow",
    "Condition": {
      "StringEquals": {
        "aws:ResourceAccount": "{{account_id}}"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToGetSecret",
    "Effect": "Allow",
    "Action": [
      "secretsmanager:GetSecretValue"
    ],
    "Resource": [
      "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[{{secret_id}}]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToDecryptSecret",
    "Effect": "Allow",
    "Action": [
      "kms:Decrypt"
    ],
    "Resource": [
      "arn:aws:kms:{{region}}:{{account_id}}:key/[{{key_id}}]"
    ],
    "Condition": {
      "StringLike": {
        "kms:ViaService": [

```

```

        "secretsmanager.*.amazonaws.com"
    ]
}
},
{
    "Sid": "AllowsAmazonQToIngestDocuments",
    "Effect": "Allow",
    "Action": [
        "qbusiness:BatchPutDocument",
        "qbusiness:BatchDeleteDocument"
    ],
    "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
},
{
    "Sid": "AllowsAmazonQToIngestPrincipalMapping",
    "Effect": "Allow",
    "Action": [
        "qbusiness:PutGroup",
        "qbusiness:CreateUser",
        "qbusiness>DeleteGroup",
        "qbusiness:UpdateUser",
        "qbusiness:ListGroups"
    ],
    "Resource": [
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}",
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}",
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}/data-source/*"
    ]
},
{
    "Sid": "AllowsAmazonQToCreateAndDeleteNI",
    "Effect": "Allow",
    "Action": [
        "ec2:CreateNetworkInterface",
        "ec2>DeleteNetworkInterface"
    ],
    "Resource": [
        "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",

```

```

        "arn:aws:ec2:{{region}}:{{account_id}}:security-group/
[[security_group]]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2:DeleteNetworkInterface"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:RequestTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
      },
      "ForAllValues:StringEquals": {
        "aws:TagKeys": [
          "AMAZON_Q"
        ]
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateTags",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateTags"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringEquals": {
        "ec2:CreateAction": "CreateNetworkInterface"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterfacePermission"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",

```

```

        "Condition": {
            "StringLike": {
                "aws:ResourceTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
            }
        },
        {
            "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
            "Effect": "Allow",
            "Action": [
                "ec2:DescribeNetworkInterfaces",
                "ec2:DescribeAvailabilityZones",
                "ec2:DescribeNetworkInterfaceAttribute",
                "ec2:DescribeVpcs",
                "ec2:DescribeRegions",
                "ec2:DescribeNetworkInterfacePermissions",
                "ec2:DescribeSubnets"
            ],
            "Resource": "*"
        }
    ]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToAssumeRoleForServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "{{source_account}}"
        },
        "ArnLike": {
          "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
        }
      }
    }
  ]
}

```

```
    }  
  }  
} ]  
}
```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Known limitations for the Amazon Q Business SharePoint (Online) connector

The Amazon Q Business SharePoint (Online) connector has the following known limitations:

- The Amazon Q SharePoint (Online) connector supports custom field mappings only for the **Files** entity.
- If an entity name has a % character in its name, the connector will skip these files due to API limitations.
- OneNote can only be crawled by the connector using a Tenant ID, and with OAuth 2.0, OAuth 2.0 refresh token, or SharePoint (Online) App Only authentication activated for SharePoint (Online) Online.
- The connector crawls the first section of a OneNote document using its default name only, even if the document is renamed.
- The connector crawls list attachments and comments only when **List Data** is also selected as an entity to be crawled.
- The connector crawls event attachments only when **Events** is also selected as an entity to be crawled.
- For SharePoint (Online) Online version, the ACL token will be in lower case. For example, if **User principal name** is *MaryMajor@domain.com* in Azure portal, the ACL token in the SharePoint Connector will be *marymajor@domain.com*. For SharePoint (Online), if the user principal name in your Azure Portal is a combination of upper case and lower case, the SharePoint API internally converts the user principal name to lower case.
- If you want to crawl nested groups using **Identity crawler**, you have to activate Local as well as AD Group Crawling.
- The User Principal Name in your Azure Portal is a combination of upper case and lower case, the SharePoint (Online) API internally converts it to lower case. Because of this, the Amazon Q SharePoint (Online) connector sets ACL in lower case.

- If you want to use **Identity Crawler** for with SharePoint (Online) to crawl nested groups, then you have to enable both Local and AD Group Crawling.

Troubleshooting your Amazon Q Business SharePoint (Online) connector

The following table provides information about error codes you may see for the Microsoft SharePoint connector and suggested troubleshooting actions.

Error code	Error message	Suggested resolution
SPE-5001	Authentication failed. Configuration might contain wrong credentials.	Provide valid credentials like username, password or client Id, client secret and tenant Id.
SPE-5002	There was a problem while connecting to Host Url and/or Domain. hostUrl and/or domain values might be incorrect .	Provide valid Host URL or Domain.
SPE-5003	Provided URL is incorrect	Provide correct URL.
SPE-5004	Inet Address validation Failed.	Provide valid Inet Address
SPE-5005	Failed : HTTP protocol violation has occurred.	Try running the connector again.
SPE-5100	There was a problem while retrieving repository Id. Repository ID might be empty or null.	Ensure that repository Id must not be null or empty.
SPE-5101	There was a problem while retrieving dataSoucelamRoleArn.	Ensure that dataSoucelamRoleArn must not be null or empty.

Error code	Error message	Suggested resolution
	Data Source IAM Role ARN might be empty or null.	
SPE-5102	There was a problem while retrieving repository configurations. Repository configurations might be empty or incorrect.	Provide valid repository configurations.
SPE-5115	There was a problem while retrieving field mapping values for event entity. Field mapping values might be empty or incorrect.	Field mapping values for event entity should be correct or non-empty.
SPE-5116	There was a problem while retrieving field mapping values for file entity. Field mapping values might be empty or incorrect.	Field mapping values for file entity should be correct or non-empty.
SPE-5117	There was a problem while retrieving field mapping values for page entity. Field mapping values might be empty or incorrect.	Field mapping values for page entity should be correct or non-empty.

Error code	Error message	Suggested resolution
SPE-5118	There was a problem while retrieving field mapping values for link entity. Field mapping values might be empty or incorrect.	Field mapping values for link entity should be correct or non-empty.
SPE-5119	There was a problem while retrieving field mapping values for comment entity. Field mapping values might be empty or incorrect.	Field mapping values for comment entity should be correct or non-empty.
SPE-5120	There was a problem while retrieving field mapping values for attachment entity. Field mapping values might be empty or incorrect.	Field mapping values for attachment entity should be correct or non-empty.
SPE-5121	There was a problem while retrieving values for crawl entities. Values might be empty or incorrect. It should be either true or false.	There might be some incorrect value given in any one of the crawling entities like – null, TRUE or any dummy string. Ensure the value must be non-empty and either true or false.
SPE-5122	There was a problem while retrieving domain. Domain might be empty or null.	Provide Client Id.

Error code	Error message	Suggested resolution
SPE-5123	There was a problem while retrieving version. Version might be empty or null.	Provide valid version and it should not be null.
SPE-5124	There was a problem while retrieving authType. Auth-Type might be empty or null.	Ensure AUTH Type in configuration must be not null.
SPE-5125	There was a problem while retrieving clientId. Client ID might be empty or null.	Provide Client Id.
SPE-5126	There was a problem while retrieving clientSecret. Client Secret might be empty or null.	Provide Client Secret.
SPE-5127	There was a problem while retrieving tenantId. Tenant ID might be empty or null.	Provide Tenant Id.
SPE-5128	There was a problem while retrieving siteUrls. Site URLs might be empty or null.	Provide at least one Site Url.
SPE-5129	There was a problem while retrieving password. Password might be empty or null.	Provide password.

Error code	Error message	Suggested resolution
SPE-5130	There was a problem while retrieving username. Username might be empty or null.	Provide username.
SPE-5131	There was a problem while retrieving username. Email was invalid.	Provide valid email address.
SPE-5132	There was a problem while retrieving url. This URL was invalid.	Provide a valid URL.
SPE-5133	There was a problem while retrieving s3CertificateName. S3 Certificate Name might be empty or null.	Ensure s3CertificateName is not null or non-empty.
SPE-5134	There was a problem while retrieving s3BucketName. S3 Bucket Name might be empty or null	Ensure s3BucketName is not null or non-empty.
SPE-5135	The provided version was not a valid Sharepoint Connector version. Version should be one of [Online, Server].	Version should be one of [Online, Server].
SPE-5136	The provided authType was not a valid Sharepoint Connector authentication method.	Provide valid authType. The value of authType should be one of [Basic, OAuth2Certificate, OAuth2].

Error code	Error message	Suggested resolution
SPE-5138	There was a problem while retrieving onPremVersion. On prem Version might be empty or null	Ensure onPremVersion is not be null or non-empty.
SPE-5139	The provided onPremVersion was not valid Sharepoint on-prem version. On prem version should be one of [2013, 2016, 2019, SubscriptionEdition].	Provide a valid onPremVersion. On prem version should be one of [2013, 2016, 2019, SubscriptionEdition].
SPE-5140	There was a problem while retrieving ldapUrl. LDAP Url might be empty or null.	Ensure ldapUrl is not null or empty.
SPE-5141	There was a problem while retrieving baseDn. Base DN might be empty or null.	Ensure baseDn is not be null or empty.
SPE-5142	There was a problem while retrieving privateKey. Private Key might be empty or null.	Please ensure privateKey is not be null or empty.
SPE-5144	There was a problem while retrieving aclConfiguration. ACL Configuration might be empty, null or invalid	Provide valid aclConfiguration. aclConfiguration should be one of [ACLWithLDAPEmailFmt, ACLWithManualEmailFmt, ACLWithUsernameFmt].

Error code	Error message	Suggested resolution
SPE-5145	There was a problem while retrieving emailDomain. Email Domain might be empty or null.	Ensure emailDomain is not null or empty.
SPE-5146	There was a problem while retrieving ldapUsername. LDAP Username might be empty or null.	Ensure ldapUser is not null or empty.
SPE-5147	There was a problem while retrieving ldapPassword. LDAP Password might be empty or null.	Ensure ldapPassword is not null or empty.
SPE-5140	SPE org ID is too large.	Org id should not be greater than 100 characters.
SPE-5141	Page name inclusion or exclusion patterns are incorrect.	Page name inclusion patterns/ Exclusion must be a list of strings.
SPE-5142	Asset name inclusion or exclusion patterns are incorrect.	Asset name inclusion patterns/ Exclusion must be a list of strings.
SPE-5143	Asset type inclusion or exclusion patterns are incorrect.	Asset type inclusion patterns/ Exclusion must be a list of strings.
SPE-5144	Invalid page root path. Please provide valid page root path.	Page path should start with /content.

Error code	Error message	Suggested resolution
SPE-5145	Invalid asset root path. Please provide valid asset root path.	Asset path should start with /content/dam.
SPE-5146	SPE page root paths list size is too large.	Page root paths list size should not be greater than 1000.
SPE-5147	SPE asset root paths list size is too large.	Asset root paths list size should not be greater than 1000.
SPE-5200	There was a problem while connecting to url:	Ensure the siteUrl exists.

Connecting SharePoint (Server) to Amazon Q Business

Microsoft SharePoint is a collaborative website building service that lets you customize web content and create web pages, web sites, document libraries, and lists. You can connect SharePoint (Server) instance to Amazon Q Business—using either the AWS Management Console or the [CreateDataSource](#) API—and create an Amazon Q web experience.

Amazon Q supports Microsoft SharePoint Server (versions 2013, 2016, 2019, and Subscription Edition).

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).
- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [SharePoint \(Server\) connector overview](#)
- [Prerequisites for connecting Amazon Q Business to SharePoint \(Server\)](#)
- [Connecting Amazon Q Business to SharePoint \(Server\) using the console](#)

- [Connecting Amazon Q Business to SharePoint \(Server\) using APIs](#)
- [How Amazon Q Business connector crawls SharePoint \(Server\) ACLs](#)
- [Amazon Q Business SharePoint \(Server\) data source connector field mappings](#)
- [IAM role for Amazon Q Business SharePoint \(Server\) connector](#)
- [Known limitations for the Amazon Q Business SharePoint \(Server\) connector](#)
- [Troubleshooting your Amazon Q Business SharePoint \(Server\) connector](#)

SharePoint (Server) connector overview

The following table gives an overview of the Amazon Q Business SharePoint (Server) connector and its supported features.

Category	Feature	Support
Security	Authentication type	NTLM, Kerberos, SharePoint App-Only
	Authentication credentials	<p>NTLM</p> <ul style="list-style-type: none"> • SharePoint (Server) admin username • SharePoint (Server) admin password <p>If you're using Email ID with Domain from IDP to crawl ACLs, then you also need to add a:</p> <ul style="list-style-type: none"> • LDAP Server Endpoint • LDAP Search Base • LDAP username • LDAP password <p>Kerberos</p> <ul style="list-style-type: none"> • SharePoint (Server) admin username • SharePoint (Server) admin password

Category	Feature	Support
		<p>If you're using Email ID with Domain from IDP to crawl ACLs, then you also need to add a:</p> <ul style="list-style-type: none"> • LDAP Server Endpoint • LDAP Search Base • LDAP username • LDAP password <p>SharePoint App-Only</p> <ul style="list-style-type: none"> • Tenant ID • SharePoint App-Only client ID • SharePoint App-Only client secret <p>If you're using Email ID with Domain from IDP to crawl ACLs, then you also need to add a:</p> <ul style="list-style-type: none"> • LDAP Server Endpoint • LDAP Search Base • LDAP username • LDAP password
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Integration with Identity Provider (IdP)	Yes. LDAP.
	Identity crawling	Yes
	VPC	Yes

Category	Feature	Support
Crawl features	Custom metadata	Yes. Supports custom metadata for File entity only.
	Entities	Yes. The following entities are supported: <ul style="list-style-type: none"> • Files • Attachments • Link • Pages • Events • Comments
	<u>Field mappings</u>	Yes. Supports both default and custom field mappings. For more information, see Field mappings .
	Filters	Yes. The following filters are supported: <ul style="list-style-type: none"> • Include/exclude by Links • Include/exclude by Pages • Include/exclude by Events • Include/exclude by file name • Include/exclude by file path • Include/exclude by file type • Include/exclude by OneNote Section name • Include/exclude by OneNote Page name
	<u>Sync mode</u>	Supports full and incremental sync.
	<u>File types</u>	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to SharePoint (Server)

The following page outlines the prerequisites you need to complete before connecting SharePoint (Server) to Amazon Q, based on the authentication mode of your choice.

Topics

- [Prerequisites for using NTLM authentication](#)
- [Prerequisites for using Kerberos authentication](#)
- [Prerequisites for using SharePoint App-Only authentication](#)

Prerequisites for using NTLM authentication

If you're using NTLM authentication, make sure you've completed the following steps in SharePoint (Server):

- Copied your SharePoint (Server) instance URLs. The format for the host URL you enter is *https://yourdomain.sharepoint.com/sites/mysite*. Your URL must start with `https` and contain `sharepoint.com`.
- Copied the domain name of your SharePoint (Server) instance URL.
- Generated an SSL certificate and uploaded it to an Amazon S3 bucket.
- Noted the username and password that you use to connect to SharePoint (Server).

(Optional) If you're using Email ID with Domain from IDP to control access to your documents, make sure you've completed the following steps:

- Copied your LDAP Server Endpoint (endpoint of LDAP server including protocol and port number). For example: *ldap://example.com:389*.
- Copied your LDAP Search Base (search base of the LDAP user). For example: *CN=Users,DC=sharepoint,DC=com*.
- Copied your LDAP user name and LDAP password.

(Optional) If using Email ID with Custom Domain for access control, complete the following step:

- Noted your custom email domain value—for example: *"amazon.com"*.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your SharePoint (Server) authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Prerequisites for using Kerberos authentication

If you're using Kerberos authentication, make sure you've completed the following steps in SharePoint (Server):

- Copied your SharePoint (Server) instance URLs. The format for the host URL you enter is *https://yourdomain.sharepoint.com/sites/mysite*. Your URL must start with `https` and contain `sharepoint.com`.
- Copied the domain name of your SharePoint (Server) instance URL.
- Generated an SSL certificate and uploaded it to an Amazon S3 bucket.
- Noted the username and password that you use to connect to SharePoint (Server).

(Optional) If you're using Email ID with Domain from IDP to control access to your documents, make sure you've completed the following steps:

- Copied your LDAP Server Endpoint (endpoint of LDAP server including protocol and port number). For example: *ldap://example.com:389*.
- Copied your LDAP Search Base (search base of the LDAP user). For example: *CN=Users,DC=sharepoint,DC=com*.
- Copied your LDAP user name and LDAP password.

(Optional) If using Email ID with Custom Domain for access control, complete the following step:

- Noted your custom email domain value—for example: *"amazon.com"*.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your SharePoint (Server) authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

 **Note**

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Prerequisites for using SharePoint App-Only authentication

If you're using SharePoint App-Only authentication, make sure you've completed the following steps in SharePoint (Server):

- Copied the SharePoint client ID generated when you registered App Only at Site Level. ClientID format is ClientID@TenantId. For example, *ffa956f3-8f89-44e7-b0e4-49670756342c@888d0b57-69f1-4fb8-957f-e1f0bedf82fe*.
- Copied the SharePoint client secret generated when you registered App Only at Site Level.

 **Important**

Note: Because client IDs and client secrets are generated for single sites only when you register SharePoint Server for App Only authentication, only one site URL is supported for SharePoint App Only authentication.

- Noted the Tenant ID of your SharePoint (Server) account.

- Noted your **LDAP Server Endpoint**, **LDAP Search Base**, **LDAP username**, and **LDAP password**.

Note

SharePoint App-Only Authentication is *not* supported for SharePoint 2013 version.

(Optional) If you're using Email ID with Domain from IDP to control access to your documents, make sure you've completed the following steps:

- Copied your LDAP Server Endpoint (endpoint of LDAP server including protocol and port number). For example: *ldap://example.com:389*.
- Copied your LDAP Search Base (search base of the LDAP user). For example: *CN=Users,DC=sharepoint,DC=com*.
- Copied your LDAP user name and LDAP password.

(Optional) If using Email ID with Custom Domain for access control, complete the following step:

- Noted your custom email domain value—for example: *"amazon.com"*.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your SharePoint (Server) authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Connecting Amazon Q Business to SharePoint (Server) using the console

The following procedure outlines how to connect Amazon Q Business to SharePoint (Server) using the AWS Management Console.

Connecting Amazon Q to SharePoint (Server)

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **SharePoint (Server)** page, enter the following information:
6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. In **Source**, enter the following information:
 - a. In **Source**, for **Hosting Method** – Choose **SharePoint Server**.
 - b. **Choose SharePoint Version** – Choose between **SharePoint 2013**, **SharePoint 2016**, **SharePoint 2019**, and **SharePoint (Subscription Edition)**.
 - c. **Site URLs specific to your SharePoint repository** – Enter the SharePoint host URLs. The format for the host URLs you enter is *https://yourcompany/sites/mysite*. The URL must start with https protocol. Separate URLs with a new line. You can add up to 100 URLs.
 - d. **Domain** – Enter the SharePoint domain.
 - e. **SSL certificate location** – Enter the Amazon S3 path to your SSL certificate file.
8. For **Web proxy – optional** – Enter the host name (without the http:// or https:// protocol), and the port number used by the host URL transport protocol. The numeric value of the port number must be between 0 and 65535.
9. For **Authorization** – You can choose to use an access control list (ACL) for controlling search results based on your end user's document access level in your SharePoint data source. Authorization using ACL is activated by default. When ACL is deactivated, no ACL information

is crawled and no access control or context filtering is available. For SharePoint Server, you can choose from the following ACL options:

- a. **Email ID with Domain from IDP** – Access control is based on email IDs that are extracted from email domains fetched from the underlying identity provider (IdP). You provide the IdP connection details in your Secrets Manager secret during **Authentication**.
- b. **Email ID with Custom Domain** – Access control is based on email IDs. Provide the email domain value. For example, "*amazon.com*". The email domain is used to construct the email ID for access control. You must enter your email domain using **Add Email Domain**.

See [Authorization](#) for more details.

10. For **Authentication**, choose between **SharePoint App-Only authentication**, **NTLM authentication**, and **Kerberos authentication**, based on your use case.

- a. Enter the following information for both **NTLM authentication** and **Kerberos authentication**:

For **AWS Secrets Manager secret** – Choose an existing secret or create a Secrets Manager secret to store your SharePoint authentication credentials. If you choose to create a secret, an AWS Secrets Manager secret window opens. Enter the following information in the window:

- **Secret name** – A name for your secret.
- **Username** – Username for your SharePoint account.
- **Password** – Password for your SharePoint account.

If using **Email ID with Domain from IDP**, also enter your:

- **LDAP Server Endpoint** – Endpoint of LDAP server, including protocol and port number. For example: *ldap://example.com:389*.
- **LDAP Search Base** – Search base of LDAP user. For example: *CN=Users,DC=sharepoint,DC=com*.
- **LDAP username** – Your LDAP username.
- **LDAP Password** – Your LDAP password.

- b. Enter the following information for **SharePoint App-Only authentication**:

For **AWS Secrets Manager secret** – Choose an existing secret or create a Secrets Manager secret to store your SharePoint authentication credentials. If you choose to create a secret, an AWS Secrets Manager secret window opens. Enter the following information in the window:

- **Secret name** – A name for your secret.
- **Client ID** – The SharePoint client ID that you generated when you registered App Only at Site Level. The ClientID format is ClientID@TenantId. For example, *ffa956f3-8f89-44e7-b0e4-49670756342c@888d0b57-69f1-4fb8-957f-e1f0bedf82fe*.
- **SharePoint client secret** – The SharePoint client secret generated when you register for App Only at Site Level.

Note: Because client IDs and client secrets are generated for single sites only when you register SharePoint Server for App Only authentication, only one site URL is supported for SharePoint App Only authentication.

If using **Email ID with Domain from IDP**, also enter your:

- **LDAP Server Endpoint** – Endpoint of LDAP server, including protocol and port number. For example: *ldap://example.com:389*.
- **LDAP Search Base** – Search base of LDAP user. For example: *CN=Users,DC=sharepoint,DC=com*.
- **LDAP username** – Your LDAP user name.
- **LDAP Password** – Your LDAP password.

11. **Configure VPC and security group – optional** – Choose whether you want to use a VPC. If you do, enter the following information:

- **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
- **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

12. **Identity crawler** – Choose to activate Amazon Q identity crawler to sync identity information. Only **Local Group Members** will be crawled by **Identity crawler**. For more information, see [Identity crawler](#).
13. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

14. In **Sync scope**, choose from the following options :
 - a. **Select entities** – Choose the entities that you want to crawl. You can select to crawl **All** entities or any combination of **Files, Attachments, Links, Pages, Events** and **List Data**.
 - b. In **Additional configuration – optional**, for **Entity regex patterns** – Add regular expression patterns for **Links, Pages, and Events** to include specific entities instead of syncing all your documents.
 - c. **Regex patterns** – Add regular expression patterns to include or exclude files by **File path, File name, File type, OneNote section name, and OneNote page name** instead of syncing all your documents. You can add up to 100 patterns.
15. In **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.
 - **Full sync** – Sync all content regardless of the previous sync status.
 - **New or modified content sync** – Sync only new and modified documents.
 - **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.

For more details, see [Sync mode](#).

16. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
17. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
18. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:

- a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
- b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

 **Note**

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

19. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

20. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

 **Note**

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to SharePoint (Server) using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the `configuration` parameter to provide a JSON schema with all other configuration information specific to your data source connector.

For an example of the API request, see [CreateDataSource](#) in the Amazon Q API Reference.

Microsoft SharePoint JSON schema

The following is the Microsoft SharePoint JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
      "type": "object",
      "properties": {
        "repositoryEndpointMetadata": {
          "type": "object",
          "properties": {
            "tenantId": {
              "type": "string",
              "pattern": "^[0-9a-f]{8}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{12}$",
              "minLength": 36,
              "maxLength": 36
            },
            "domain": {
              "type": "string"
            },
            "siteUrls": {
              "type": "array",
              "items": {
                "type": "string",
                "pattern": "https://.*"
              }
            }
          }
        },
        "repositoryAdditionalProperties": {
          "type": "object",
          "properties": {
            "s3bucketName": {
              "type": "string"
            },
            "s3certificateName": {
              "type": "string"
            },
            "authType": {
              "type": "string",
              "enum": [
                "OAuth2",
                "OAuth2Certificate",

```

```

    "OAuth2App",
    "OAuth2_RefreshToken",
    "Basic",
    "NTLM",
    "Kerberos"
  ]
},
"version": {
  "type": "string",
  "enum": [
    "Server",
    "Online"
  ]
},
"onPremVersion": {
  "type": "string",
  "enum": [
    "",
    "2013",
    "2016",
    "2019",
    "SubscriptionEdition"
  ]
}
},
"required": [
  "authType",
  "version"
]
}
},
"required": [
  "siteUrls",
  "domain",
  "repositoryAdditionalProperties"
]
}
},
"required": [
  "repositoryEndpointMetadata"
]
},
"repositoryConfigurations": {
  "type": "object",

```

```

"properties": {
  "event": {
    "type": "object",
    "properties": {
      "fieldMappings": {
        "type": "array",
        "items": [
          {
            "type": "object",
            "properties": {
              "indexFieldName": {
                "type": "string"
              },
              "indexFieldType": {
                "type": "string",
                "enum": [
                  "STRING",
                  "STRING_LIST",
                  "DATE"
                ]
              },
              "dataSourceFieldName": {
                "type": "string"
              },
              "dateFieldFormat": {
                "type": "string",
                "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
              }
            },
            "required": [
              "indexFieldName",
              "indexFieldType",
              "dataSourceFieldName"
            ]
          }
        ]
      },
      "required": [
        "fieldMappings"
      ]
    },
    "page": {
      "type": "object",

```

```

"properties": {
  "fieldMappings": {
    "type": "array",
    "items": [
      {
        "type": "object",
        "properties": {
          "indexFieldName": {
            "type": "string"
          },
          "indexFieldType": {
            "type": "string",
            "enum": [
              "STRING",
              "DATE",
              "LONG"
            ]
          },
          "dataSourceFieldName": {
            "type": "string"
          },
          "dateFieldFormat": {
            "type": "string",
            "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
          }
        },
        "required": [
          "indexFieldName",
          "indexFieldType",
          "dataSourceFieldName"
        ]
      }
    ]
  },
  "required": [
    "fieldMappings"
  ]
},
"file": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",

```



```

"items": [
  {
    "type": "object",
    "properties": {
      "indexFieldName": {
        "type": "string"
      },
      "indexFieldType": {
        "type": "string",
        "enum": [
          "STRING",
          "DATE",
          "LONG"
        ]
      },
      "dataSourceFieldName": {
        "type": "string"
      },
      "dateFieldFormat": {
        "type": "string",
        "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
      }
    },
    "required": [
      "indexFieldName",
      "indexFieldType",
      "dataSourceFieldName"
    ]
  }
],
"required": [
  "fieldMappings"
],
"link": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": [
        {
          "type": "object",

```

```
"properties": {
  "indexFieldName": {
    "type": "string"
  },
  "indexFieldType": {
    "type": "string",
    "enum": [
      "STRING",
      "STRING_LIST",
      "DATE"
    ]
  },
  "dataSourceFieldName": {
    "type": "string"
  },
  "dateFieldFormat": {
    "type": "string",
    "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
  }
},
"required": [
  "indexFieldName",
  "indexFieldType",
  "dataSourceFieldName"
]
},
"required": [
  "fieldMappings"
]
},
"attachment": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": [
        {
          "type": "object",
          "properties": {
            "indexFieldName": {
              "type": "string"
            }
          }
        }
      ]
    }
  }
}
```

```

    },
    "indexFieldType": {
      "type": "string",
      "enum": [
        "STRING",
        "STRING_LIST",
        "DATE"
      ]
    },
    "dataSourceFieldName": {
      "type": "string"
    },
    "dateFieldFormat": {
      "type": "string",
      "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
    }
  ],
  "required": [
    "indexFieldName",
    "indexFieldType",
    "dataSourceFieldName"
  ]
}
]
}
},
"required": [
  "fieldMappings"
]
},
"comment": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": [
        {
          "type": "object",
          "properties": {
            "indexFieldName": {
              "type": "string"
            },
            "indexFieldType": {
              "type": "string",

```

```

    "enum": [
      "STRING",
      "STRING_LIST",
      "DATE"
    ]
  },
  "dataSourceFieldName": {
    "type": "string"
  },
  "dateFieldFormat": {
    "type": "string",
    "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
  }
},
"required": [
  "indexFieldName",
  "indexFieldType",
  "dataSourceFieldName"
]
}
]
}
},
"required": [
  "fieldMappings"
]
}
}
},
"additionalProperties": {
  "type": "object",
  "properties": {
    "eventTitleFilterRegEx": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    "pageTitleFilterRegEx": {
      "type": "array",
      "items": {
        "type": "string"
      }
    }
  }
},
},

```

```
"linkTitleFilterRegex": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"inclusionFilePath": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"exclusionFilePath": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"inclusionFileTypePatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"exclusionFileTypePatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"inclusionFileNamePatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"exclusionFileNamePatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"inclusionOneNoteSectionNamePatterns": {
  "type": "array",
```

```
"items": {
  "type": "string"
},
"exclusionOneNoteSectionNamePatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"inclusionOneNotePageNamePatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"exclusionOneNotePageNamePatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"crawlFiles": {
  "type": "boolean"
},
"crawlPages": {
  "type": "boolean"
},
"crawlEvents": {
  "type": "boolean"
},
"crawlComments": {
  "type": "boolean"
},
"crawlLinks": {
  "type": "boolean"
},
"crawlAttachments": {
  "type": "boolean"
},
"crawlListData": {
  "type": "boolean"
},
"crawlAcl": {
```

```
    "type": "boolean"
  },
  "aclConfiguration": {
    "type": "string",
    "enum": [
      "ACLWithLDAPEmailFmt",
      "ACLWithManualEmailFmt",
      "ACLWithUsernameFmt"
    ]
  },
  "emailDomain": {
    "type": "string"
  },
  "isCrawlLocalGroupMapping": {
    "type": "boolean"
  },
  "isCrawlAdGroupMapping": {
    "type": "boolean"
  },
  "proxyHost": {
    "type": "string"
  },
  "proxyPort": {
    "type": "string"
  }
},
"required": [
]
},
"type": {
  "type": "string",
  "pattern": "SHAREPOINTV2"
},
"enableIdentityCrawler": {
  "type": "boolean"
},
"syncMode": {
  "type": "string",
  "enum": [
    "FULL_CRAWL",
    "FORCED_FULL_CRAWL",
    "CHANGE_LOG"
  ]
},
},
```

```

"secretArn": {
  "type": "string",
  "minLength": 20,
  "maxLength": 2048
},
"version": {
  "type": "string",
  "anyOf": [
    {
      "pattern": "1.0.0"
    }
  ]
},
"required": [
  "connectionConfiguration",
  "repositoryConfigurations",
  "enableIdentityCrawler",
  "syncMode",
  "additionalProperties",
  "secretArn",
  "type"
]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.
repositoryEndpointMetadata	The endpoint information for the data source.
tenantId	The tenant id of your SharePoint account.
domain	The domain of your SharePoint account.
siteUrls	The host URLs of your SharePoint account.
repositoryAdditionalProperties	Additional properties to connect with your repository endpoint.

Configuration	Description
s3bucketName	The name of the Amazon S3 bucket that stores your Azure AD self-signed X.509 certificate.
s3certificateName	The name of the SSL certificate stored in your Amazon S3 bucket.
authType	The type of authentication you are using: OAuth2, OAuth2Certificate , OAuth2App , or Basic.
version	The SharePoint version you are using: Online.
onPremVersion	Not required if you are using SharePoint Online.
repositoryConfigurations <ul style="list-style-type: none"> • event • page • file • link • attachment • comment 	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings. <p>A list of objects that map the attributes or field names of your SharePoint (Server) pages and assets to Amazon Q index field names.</p>
additionalProperties	Additional configuration options for your content in your data source.

Configuration	Description
<ul style="list-style-type: none">• eventTitleFilterRegEx• pageTitleFilterRegEx• linkTitleFilterRegEx• inclusionFilePath• exclusionFilePath• inclusionFileTypePatterns• exclusionFileTypePatterns• inclusionFileNamePatterns• exclusionFileNamePatterns• inclusionOneNoteSectionNamePatterns• exclusionOneNoteSectionNamePatterns• inclusionOneNotePageNamePatterns• exclusionOneNotePageNamePatterns• aclConfiguration• emailDomain• proxyHost• proxyPort	<p>A list of regular expression patterns to include/exclude specific files in your SharePoint data source. Files that match the patterns are included in the index. File that don't match the patterns are excluded from the index. If a file matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence, and the file isn't included in the index.</p>

Configuration	Description
<ul style="list-style-type: none"> • <code>crawlFiles</code> • <code>crawlPages</code> • <code>crawlEvents</code> • <code>crawlComments</code> • <code>crawlLinks</code> • <code>crawlAttachments</code> • <code>crawlListData</code> • <code>crawlAcl</code> • <code>isCrawlLocalGroupMapping</code> • <code>isCrawlAdGroupMapping</code> 	<p>Input TRUE to index.</p>
<p><code>type</code></p>	<p>Specify <code>SHAREPOINTV2</code> as your data source type</p>
<p><code>enableIdentityCrawler</code></p>	<p>true to activate identity crawler. Identity crawler is activated by default. Crawling identity information on users and groups with access to specific documents is useful for user context filtering. Search results are filtered based on the user or their group access to documents. See Identity crawler for more information.</p>

Configuration	Description
syncMode	<p>Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose between the following options:</p> <ul style="list-style-type: none"> • Use <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index • Use <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index • Use <code>CHANGE_LOG</code> to incrementally crawl only new and modified content each time your data source syncs with your index
secretARN	<p>The Amazon Resource Name (ARN) of an AWS Secrets Manager secret that contains the key-value pairs required to connect to your SharePoint. If you use basic authentication provide the username and password. If you use OAuth 2.0 authentication, provide the username, password, client ID, and client secret.</p>
version	<p>The version of this template that's currently supported.</p>

How Amazon Q Business connector crawls SharePoint (Server) ACLs

When you connect an SharePoint (Server) data source to Amazon Q Business, Amazon Q Business crawls ACL information attached to a document (user and group information) from your SharePoint (Server) instance. If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

To filter using a username, use the **User principal name** from your Azure portal. For example, johnstiles@kendra.onmicrosoft.com.

When you use a SharePoint group for user context filtering, calculate the group ID as follows:

For local groups

1. Get the site name. For example, `https://host.onmicrosoft.com/sites/siteName`.
2. Take the SHA256 hash of the site name. For example, `430a6b90503eef95c89295c8999c7981`.
3. Create the group ID by concatenating the SHA256 hash with a vertical bar (|) and the group name. For example, if the group name is "local group name", the group ID is the following:

`"430a6b90503eef95c89295c8999c7981 | localGroupName"` (with a space before and after the vertical bar).

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business SharePoint (Server) data source connector field mappings

To help you structure data for retrieval and chat filtering, Amazon Q Business crawls data source document attributes or metadata and maps them to fields in your Amazon Q index.

Amazon Q has reserved fields that it uses when querying your application. When possible, Amazon Q automatically maps these built-in fields to attributes in your data source. If a built-in field doesn't have a default mapping, or if you want to map additional index fields, use the custom field mappings to specify how a data source attribute maps to your Amazon Q application. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

⚠ Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q Sharepoint connector supports the following entities and the associated reserved and custom attributes.

ℹ Note

You can map any Sharepoint field to the document title or document body Amazon Q reserved/default index fields.

Supported entities and field mappings

- [Files](#)
- [Events](#)
- [Pages](#)
- [Links](#)
- [Attachments](#)
- [Comments](#)

Files

Sharepoint field name	Index field name	Description	Data type
title	sp_title	Custom	String
sourceUri	_source_uri	Default	String
checkInComment	sp_checkInComment	Custom	String
size	sp_sizeLong	Custom	Long (numeric)

Sharepoint field name	Index field name	Description	Data type
lastModifiedDateTime	_last_updated_at	Default	Date
createdAt	_created_at	Default	Date
author	_authors	Default	String list
majorVersion	sp_majorVersion	Custom	String
uiVersionLabel	sp_uiVersionLabel	Custom	String
uniqueId	sp_uniqueId	Custom	String
irmEnabled	sp_irmEnabled	Custom	String
checkOutType	sp_checkOutType	Custom	String
category	_category	Default	String
modifiedBy	sp_modifiedBy	Custom	String
level	sp_level	Custom	String
uiVersion	sp_uiVersion	Custom	String
contentTag	sp_contentTag	Custom	String
eTag	sp_eTag	Custom	String
oneNoteDocument	sp_oneNoteDocument	Custom	String
oneNoteSection	sp_oneNoteSection	Custom	String
oneNotePage	sp_oneNotePage	Custom	String

Events

Sharepoint field name	Index field name	Description	Data type
title	sp_title	Custom	String
lastModifiedDateTime	_last_updated_at	Default	Date
sourceUri	_source_uri	Default	String
attachments	sp_hasAttachments	Custom	String
createdDate	_created_at	Default	Date
authorId	sp_authorId	Custom	String
editorId	sp_editorId	Custom	String
location	sp_location	Custom	String
eventDate	sp_eventDate	Custom	Date
eventEndDate	sp_eventEndDate	Custom	Date
ifRecurrence	sp_ifRecurrence	Custom	String
ifAllDayEvent	sp_ifAllDayEvent	Custom	String
category	_category	Default	String
eventCategory	sp_eventcategory	Custom	String

Pages

Sharepoint field name	Index field name	Description	Data type
createdDateTime	_created_at	Default	Date

Sharepoint field name	Index field name	Description	Data type
lastModifiedDateTime	_last_updated_at	Default	Date
title	sp_title	Custom	String
sourceUri	_source_uri	Default	String
firstPublishedDate	sp_firstPublishedDate	Custom	Date
authorId	sp_authorId	Custom	String
editorId	sp_editorId	Custom	String
category	_category	Default	String

Links

Sharepoint field name	Index field name	Description	Data type
createdAt	_created_at	Default	Date
lastModifiedDateTime	_last_updated_at	Default	Date
title	sp_title	Custom	String
sourceUri	_source_uri	Default	String
fileType	sp_fileType	Custom	String
fileDirPath	sp_fileDirPath	Custom	String
firstPublishedDate	sp_firstPublishedDate	Custom	Date
authorId	sp_authorId	Custom	String

Sharepoint field name	Index field name	Description	Data type
editorId	sp_editorId	Custom	String
category	_category	Default	String
size	sp_sizeLong	Custom	Long (numeric)

Attachments

Sharepoint field name	Index field name	Description	Data type
title	sp__title	Custom	String
parentCreatedDate	_created_at	Default	Date
sourceUri	_source_uri	Default	String
parentModifiedDate	_last_updated_at	Custom	Date
parentListId	sp_parentListId	Custom	String
parentTitle	sp_parentTitle	Custom	String
category	_category	Default	String

Comments

Sharepoint field name	Index field name	Description	Data type
createdDateTime	_created_at	Default	Date
likedBy	sp_likedBy	Custom	String
sourceUri	_source_uri	Custom	String

Sharepoint field name	Index field name	Description	Data type
isReply	sp_isReply	Custom	String
author	_authors	Default	String list
listId	sp_listId	Custom	String
category	_category	Default	String
replyCount	sp_replyCount	Custom	String
parentTitle	sp_parentTitle	Custom	String

IAM role for Amazon Q Business SharePoint (Server) connector

Note

(Optional) If you use **Azure App-Only authentication**, you also need to add permissions for Amazon Q to access the certificate stored in your Amazon S3 bucket.

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the BatchPutDocument and BatchDeleteDocument operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.

- Permission to access the SSL certificate stored in your Amazon S3 bucket.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AllowsAmazonQToGetS3Objects",
    "Action": [
      "s3:GetObject"
    ],
    "Resource": [
      "arn:aws:s3:::{{input_bucket_name}}/*"
    ],
    "Effect": "Allow",
    "Condition": {
      "StringEquals": {
        "aws:ResourceAccount": "{{account_id}}"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToGetSecret",
    "Effect": "Allow",
    "Action": [
      "secretsmanager:GetSecretValue"
    ],
    "Resource": [
      "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToDecryptSecret",
    "Effect": "Allow",
    "Action": [
      "kms:Decrypt"
    ],
    "Resource": [
      "arn:aws:kms:{{region}}:{{account_id}}:key/[key_id]"
    ],
    "Condition": {
      "StringLike": {
        "kms:ViaService": [

```

```

        "secretsmanager.*.amazonaws.com"
    ]
}
},
{
    "Sid": "AllowsAmazonQToIngestDocuments",
    "Effect": "Allow",
    "Action": [
        "qbusiness:BatchPutDocument",
        "qbusiness:BatchDeleteDocument"
    ],
    "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
},
{
    "Sid": "AllowsAmazonQToIngestPrincipalMapping",
    "Effect": "Allow",
    "Action": [
        "qbusiness:PutGroup",
        "qbusiness:CreateUser",
        "qbusiness>DeleteGroup",
        "qbusiness:UpdateUser",
        "qbusiness:ListGroups"
    ],
    "Resource": [
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}",
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}",
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}/data-source/*"
    ]
},
{
    "Sid": "AllowsAmazonQToCreateAndDeleteNI",
    "Effect": "Allow",
    "Action": [
        "ec2:CreateNetworkInterface",
        "ec2>DeleteNetworkInterface"
    ],
    "Resource": [
        "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",

```

```

        "arn:aws:ec2:{{region}}:{{account_id}}:security-group/
[[security_group]]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2:DeleteNetworkInterface"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:RequestTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
      },
      "ForAllValues:StringEquals": {
        "aws:TagKeys": [
          "AMAZON_Q"
        ]
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateTags",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateTags"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringEquals": {
        "ec2:CreateAction": "CreateNetworkInterface"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterfacePermission"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",

```

```

        "Condition": {
            "StringLike": {
                "aws:ResourceTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
            }
        },
        {
            "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
            "Effect": "Allow",
            "Action": [
                "ec2:DescribeNetworkInterfaces",
                "ec2:DescribeAvailabilityZones",
                "ec2:DescribeNetworkInterfaceAttribute",
                "ec2:DescribeVpcs",
                "ec2:DescribeRegions",
                "ec2:DescribeNetworkInterfacePermissions",
                "ec2:DescribeSubnets"
            ],
            "Resource": "*"
        }
    ]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToAssumeRoleForServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "{{source_account}}"
        },
        "ArnLike": {
          "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
        }
      }
    }
  ]
}

```

```
    }  
  }  
} ]  
}
```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Known limitations for the Amazon Q Business SharePoint (Server) connector

The Amazon Q Business SharePoint (Server) connector has the following known limitations:

- The Amazon Q SharePoint (Server) connector supports custom field mappings only for the **Files** entity.
- For all SharePoint (Server) Server versions, the ACL token must be in lower case. For **Email with Domain from IDP** and **Email ID with Custom Domain** ACL, for example: *user@sharepoint2019.com*. For **Domain\User with Domain** ACL, for example: *sharepoint2013\user*.
- The connector does not support change log mode/**New or modified content sync** for SharePoint (Server) 2013.
- If an entity name has a % character in its name, the connector will skip these files due to API limitations.
- OneNote can only be crawled by the connector using a Tenant ID, and with OAuth 2.0, OAuth 2.0 refresh token, or SharePoint (Server) App Only authentication activated for SharePoint (Server) Online.
- The connector crawls the first section of a OneNote document using its default name only, even if the document is renamed.
- The connector crawls links in SharePoint (Server) 2019, SharePoint (Server) Online, and Subscription Edition, only if **Pages** and **Files** are selected as entities to be crawled in addition to **Links**.
- The connector crawls links in SharePoint (Server) 2013 and SharePoint (Server) 2016 if **Links** is selected as an entity to be crawled.
- The connector crawls list attachments and comments only when **List Data** is also selected as an entity to be crawled.
- The connector crawls event attachments only when **Events** is also selected as an entity to be crawled.

- If you want to crawl nested groups using **Identity crawler**, you have to activate Local as well as AD Group Crawling.
- If you want to use **Identity Crawler** for with SharePoint (Server) to crawl nested groups, then you have to enable both Local and AD Group Crawling.

Troubleshooting your Amazon Q Business SharePoint (Server) connector

The following table provides information about error codes you may see for the Microsoft SharePoint connector and suggested troubleshooting actions.

Error code	Error message	Suggested resolution
SPE-5001	Authentication failed. Configuration might contain wrong credentials.	Provide valid credentials like username, password or client Id, client secret and tenant Id.
SPE-5002	There was a problem while connecting to Host Url and/or Domain. hostUrl and/or domain values might be incorrect .	Provide valid Host URL or Domain.
SPE-5003	Provided URL is incorrect	Provide correct URL.
SPE-5004	Inet Address validation Failed.	Provide valid Inet Address
SPE-5005	Failed : HTTP protocol violation has occurred.	Try running the connector again.
SPE-5100	There was a problem while retrieving repository Id. Repository ID might be empty or null.	Ensure that repository Id must not be null or empty.

Error code	Error message	Suggested resolution
SPE-5101	There was a problem while retrieving dataSoucelamRoleArn. Data Source IAM Role ARN might be empty or null.	Ensure that dataSoucelamRoleArn must not be null or empty.
SPE-5102	There was a problem while retrieving repository configurations. Repository configurations might be empty or incorrect.	Provide valid repository configurations.
SPE-5115	There was a problem while retrieving field mapping values for event entity. Field mapping values might be empty or incorrect.	Field mapping values for event entity should be correct or non-empty.
SPE-5116	There was a problem while retrieving field mapping values for file entity. Field mapping values might be empty or incorrect.	Field mapping values for file entity should be correct or non-empty.
SPE-5117	There was a problem while retrieving field mapping values for page entity. Field mapping values might be empty or incorrect.	Field mapping values for page entity should be correct or non-empty.

Error code	Error message	Suggested resolution
SPE-5118	There was a problem while retrieving field mapping values for link entity. Field mapping values might be empty or incorrect.	Field mapping values for link entity should be correct or non-empty.
SPE-5119	There was a problem while retrieving field mapping values for comment entity. Field mapping values might be empty or incorrect.	Field mapping values for comment entity should be correct or non-empty.
SPE-5120	There was a problem while retrieving field mapping values for attachment entity. Field mapping values might be empty or incorrect.	Field mapping values for attachment entity should be correct or non-empty.
SPE-5121	There was a problem while retrieving values for crawl entities. Values might be empty or incorrect. It should be either true or false.	There might be some incorrect value given in any one of the crawling entities like – null, TRUE or any dummy string. Ensure the value must be non-empty and either true or false.
SPE-5122	There was a problem while retrieving domain. Domain might be empty or null.	Provide Client Id.

Error code	Error message	Suggested resolution
SPE-5123	There was a problem while retrieving version. Version might be empty or null.	Provide valid version and it should not be null.
SPE-5124	There was a problem while retrieving authType. Auth-Type might be empty or null.	Ensure AUTH Type in configuration must be not null.
SPE-5125	There was a problem while retrieving clientId. Client ID might be empty or null.	Provide Client Id.
SPE-5126	There was a problem while retrieving clientSecret. Client Secret might be empty or null.	Provide Client Secret.
SPE-5127	There was a problem while retrieving tenantId. Tenant ID might be empty or null.	Provide Tenant Id.
SPE-5128	There was a problem while retrieving siteUrls. Site URLs might be empty or null.	Provide at least one Site Url.
SPE-5129	There was a problem while retrieving password. Password might be empty or null.	Provide password.

Error code	Error message	Suggested resolution
SPE-5130	There was a problem while retrieving username. Username might be empty or null.	Provide username.
SPE-5131	There was a problem while retrieving username. Email was invalid.	Provide valid email address.
SPE-5132	There was a problem while retrieving url. This URL was invalid.	Provide a valid URL.
SPE-5133	There was a problem while retrieving s3CertificateName. S3 Certificate Name might be empty or null.	Ensure s3CertificateName is not null or non-empty.
SPE-5134	There was a problem while retrieving s3BucketName. S3 Bucket Name might be empty or null	Ensure s3BucketName is not null or non-empty.
SPE-5135	The provided version was not a valid Sharepoint Connector version. Version should be one of [Online, Server].	Version should be one of [Online, Server].
SPE-5136	The provided authType was not a valid Sharepoint Connector authentication method.	Provide valid authType. The value of authType should be one of [Basic, OAuth2Certificate, OAuth2].

Error code	Error message	Suggested resolution
SPE-5138	There was a problem while retrieving onPremVersion. On prem Version might be empty or null	Ensure onPremVersion is not be null or non-empty.
SPE-5139	The provided onPremVersion was not valid Sharepoint on-prem version. On prem version should be one of [2013, 2016, 2019, SubscriptionEdition].	Provide a valid onPremVersion. On prem version should be one of [2013, 2016, 2019, SubscriptionEdition].
SPE-5140	There was a problem while retrieving ldapUrl. LDAP Url might be empty or null.	Ensure ldapUrl is not null or empty.
SPE-5141	There was a problem while retrieving baseDn. Base DN might be empty or null.	Ensure baseDn is not be null or empty.
SPE-5142	There was a problem while retrieving privateKey. Private Key might be empty or null.	Please ensure privateKey is not be null or empty.
SPE-5144	There was a problem while retrieving aclConfiguration. ACL Configuration might be empty, null or invalid	Provide valid aclConfiguration. aclConfiguration should be one of [ACLWithLDAPEmailFmt, ACLWithManualEmailFmt, ACLWithUsernameFmt].

Error code	Error message	Suggested resolution
SPE-5145	There was a problem while retrieving emailDomain. Email Domain might be empty or null.	Ensure emailDomain is not null or empty.
SPE-5146	There was a problem while retrieving ldapUsername. LDAP Username might be empty or null.	Ensure ldapUser is not null or empty.
SPE-5147	There was a problem while retrieving ldapPassword. LDAP Password might be empty or null.	Ensure ldapPassword is not null or empty.
SPE-5140	SPE org ID is too large.	Org id should not be greater than 100 characters.
SPE-5141	Page name inclusion or exclusion patterns are incorrect.	Page name inclusion patterns/ Exclusion must be a list of strings.
SPE-5142	Asset name inclusion or exclusion patterns are incorrect.	Asset name inclusion patterns/ Exclusion must be a list of strings.
SPE-5143	Asset type inclusion or exclusion patterns are incorrect.	Asset type inclusion patterns/ Exclusion must be a list of strings.
SPE-5144	Invalid page root path. Please provide valid page root path.	Page path should start with /content.

Error code	Error message	Suggested resolution
SPE-5145	Invalid asset root path. Please provide valid asset root path.	Asset path should start with /content/dam.
SPE-5146	SPE page root paths list size is too large.	Page root paths list size should not be greater than 1000.
SPE-5147	SPE asset root paths list size is too large.	Asset root paths list size should not be greater than 1000.
SPE-5200	There was a problem while connecting to url:	Ensure the siteUrl exists.

Connecting Microsoft SQL Server to Amazon Q Business

Microsoft SQL Server is an relational database management system (RDBMS) developed by Microsoft. You can connect your Microsoft SQL Server instance to Amazon Q Business—using either the AWS Management Console, CLI, or the [CreateDataSource](#) API—and create an Amazon Q web experience.

The Amazon Q Microsoft SQL Server data source connector supports MS SQL Server 2019.

Important

As a best practice, provide Amazon Q with read-only database credentials. Also, avoid adding tables with sensitive data or personal identifiable information (PII).

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).
- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [Microsoft SQL Server connector overview](#)
- [Prerequisites for connecting Amazon Q Business to Microsoft SQL Server](#)
- [Connecting Amazon Q Business to Microsoft SQL Server using the console](#)
- [Connecting Amazon Q Business to Microsoft SQL Server using APIs](#)
- [How Amazon Q Business connector crawls Microsoft SQL Server ACLs](#)
- [Amazon Q Business Microsoft SQL Server data source connector field mappings](#)
- [IAM role for Amazon Q Business Microsoft SQL Server connector](#)
- [Known limitations for the Amazon Q Business Microsoft SQL Server connector](#)

Microsoft SQL Server connector overview

The following table gives an overview of the Amazon Q Business Microsoft SQL Server connector and its supported features.

Category	Feature	Support
Security	Authentication type	Basic
	Authentication credentials	<ul style="list-style-type: none"> • Username of database user • Password of database user
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Driver version	Microsoft SQL Server – 10.2.0.jre11
	Data source version	SQL Server 2019
	Identity crawling	No
	VPC	Yes
Crawl features	Custom metadata	Yes
	Entities	Yes. The following entities are supported: <ul style="list-style-type: none"> • Document

Category	Feature	Support
		<div style="border: 1px solid #add8e6; border-radius: 10px; padding: 10px;"> <p>Note</p> <p>Each database row is considered an individual searchable Amazon Q document.</p> </div>
	Field mappings	Yes. Supports both default and custom field mappings. For more information, see Field mappings .
	Sync mode	Supports full and incremental sync.
	File types	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to Microsoft SQL Server

Before you begin, make sure that you have completed the following prerequisites.

In Microsoft SQL Server, make sure you have:

- Noted your database user name and password.

Important

As a best practice, provide Amazon Q with read-only database credentials.

- Copied your database host URL, port, and instance.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your Microsoft SQL Server authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Connecting Amazon Q Business to Microsoft SQL Server using the console

The following procedure outlines how to connect Amazon Q Business to Microsoft SQL Server using the AWS Management Console.

Connecting Amazon Q to Microsoft SQL Server

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **Microsoft SQL Server** page, enter the following information:
6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. In **Source**, enter the following information:
 - a. **Host** – Enter the database host name.
 - b. **Port** – Enter the database port.
 - c. **Instance** – Enter the database instance.
 - d. **Enable SSL certificate location** – Choose to enter the Amazon S3 path to your SSL certificate file.

8. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

 **Note**

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

9. In **Authentication** – Enter the following information for your **AWS Secrets Manager secret**.
 - a. **Secret name** – A name for your secret.
 - b. For **Database user name**, and **Password** – Enter the authentication credential values you copied from your database.
 - c. Choose **Save**.
10. **Configure VPC and security group – optional** – Choose whether you want to use a VPC. If you do, enter the following information:
 - a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
 - b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

11. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

12. In **Sync scope**, enter the following information:
 - **SQL query** – Enter SQL query statements like SELECT and JOIN operations. SQL queries must be less than 1000 characters and not contain any semi-colons (;). Amazon Q will crawl all database content that matches your query.

If a table name has special characters, put it in square brackets "[]" in the SQL query. For example: `select * from [my-database-table]`.

- **Primary key column** – Provide the primary key for the database table. This identifies a table within your database.
- **Title column** – Provide the name of the document title column within your database table.
- **Body column** – Provide the name of the document body column within your database table.

13. In **Additional configuration** – *optional* – Configure the following settings:

- **Change-detecting columns** – Enter the names of the columns that Amazon Q will use to detect content changes. Amazon Q will re-index content when there is a change in any of these columns.
- **Users' IDs column** – Enter the name of the column which contains User IDs to be allowed access to content.
- **Groups column** – Enter the name of the column that contains groups to be allowed access to content.
- **Source URLs column** – Enter the name of the column which contains Source URLs to be indexed.
- **Time stamps column** – Enter the name of the column which contains time stamps. Amazon Q uses time stamp information to detect changes in your content and sync only changed content.
- **Time zones column** – Enter the name of the column which contains time zones for the content to be crawled.
- **Time stamps format** – Enter the name of the column which contains time stamp formats to use to detect content changes and re-sync your content.

14. In **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.

- **Full sync** – Sync all content regardless of the previous sync status.
- **New or modified content sync** – Sync only new and modified documents.
- **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.

For more details, see [Sync mode](#).

15. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
16. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
17. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:
 - a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
 - b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

 **Note**

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

18. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

19. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

 **Note**

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to Microsoft SQL Server using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the configuration parameter to provide a JSON schema with all other configuration information specific to your data source connector.

Microsoft SQL Server JSON schema

The following is the Microsoft SQL Server JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
      "type": "object",
      "properties": {
        "repositoryEndpointMetadata": {
          "type": "object",
          "properties": {
            "dbType": {
              "type": "string",
              "enum": [
                "mysql",
                "db2",
                "postgresql",
                "oracle",
                "sqlserver"
              ]
            },
            "dbHost": {
              "type": "string"
            },
            "dbPort": {
              "type": "string"
            },
            "dbInstance": {
              "type": "string"
            }
          }
        },
        "required": [
          "dbType",
          "dbHost",
```

```

        "dbPort",
        "dbInstance"
    ]
}
},
"required": [
    "repositoryEndpointMetadata"
]
},
"repositoryConfigurations": {
    "type": "object",
    "properties": {
        "document": {
            "type": "object",
            "properties": {
                "fieldMappings": {
                    "type": "array",
                    "items": [
                        {
                            "type": "object",
                            "properties": {
                                "indexFieldName": {
                                    "type": "string"
                                },
                                "indexFieldType": {
                                    "type": "string"
                                },
                                "dataSourceFieldName": {
                                    "type": "string"
                                }
                            }
                        },
                    ]
                },
                "required": [
                    "indexFieldName",
                    "indexFieldType",
                    "dataSourceFieldName"
                ]
            }
        }
    ],
    "required": [
        "fieldMappings"
    ]
}
}

```



```
    },
    "required": [
    ]
  },
  "additionalProperties": {
    "type": "object",
    "properties": {
      "primaryKey": {
        "type": "string"
      },
      "titleColumn": {
        "type": "string"
      },
      "bodyColumn": {
        "type": "string"
      },
      "sqlQuery": {
        "type": "string",
        "not": {
          "pattern": ";+"
        }
      },
      "timestampColumn": {
        "type": "string"
      },
      "timestampFormat": {
        "type": "string"
      },
      "timezone": {
        "type": "string"
      },
      "changeDetectingColumns": {
        "type": "array",
        "items": {
          "type": "string"
        }
      },
      "allowedUsersColumn": {
        "type": "string"
      },
      "allowedGroupsColumn": {
        "type": "string"
      },
      "sourceURIColumn": {
```

```

        "type": "string"
    },
    "serverlessAurora": {
        "type": "string",
        "enum": ["true", "false"]
    }
},
"required": ["primaryKey", "titleColumn", "bodyColumn", "sqlQuery"]
},
"type" : {
    "type" : "string",
    "pattern": "JDBC"
},
"syncMode": {
    "type": "string",
    "enum": [
        "FORCED_FULL_CRAWL",
        "FULL_CRAWL"
    ]
},
"secretArn": {
    "type": "string"
}
},
"version": {
    "type": "string",
    "anyOf": [
        {
            "pattern": "1.0.0"
        }
    ]
}
},
"required": [
    "connectionConfiguration",
    "repositoryConfigurations",
    "syncMode",
    "additionalProperties",
    "secretArn",
    "type"
]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.
repositoryEndpointMetadata	<p>Required configuration information for connecting your data source.</p> <ul style="list-style-type: none"> • dbType—The type of Java database you are using, whether <code>mysql</code>, <code>db2</code>, <code>postgresql</code>, <code>oracle</code>, or <code>sqlserver</code>. • dbHost—The database host name. • dbPort—The database port. • dbInstance—The database instance.
repositoryConfigurations	<p>Configuration information for the content of the data source. For example, configuring specific types of content and field mappings. Specify the type of data source and the secret ARN.</p>
document	<p>A list of objects that map the attributes or field names of your database content to Amazon Q index field names. For more information, see Mapping data source fields.</p>
additionalProperties	<p>Additional configuration options for your content in your data source. Use to include or exclude specific content in your database data source.</p>
primaryKey	<p>Provide the primary key for the database table. This identifies a table within your database.</p>
titleColumn	<p>Provide the name of the document title column within your database table.</p>

Configuration	Description
bodyColumn	Provide the name of the document title column within your database table.
sqlQuery	Enter SQL query statements like SELECT and JOIN operations. SQL queries must be less than 1000 characters and not contain any semi-colons (;). Amazon Q will crawl all database content that matches your query. If a table name has special characters, put it in square brackets "[]" in the SQL query. For example: <code>select * from [my-database-table] .</code>
timestampColumn	Enter the name of the column which contains time stamps. Amazon Q uses time stamp information to detect changes in your content and sync only changed content.
timestampFormat	Enter the name of the column which contains time stamp formats to use to detect content changes and re-sync your content.
timezone	Enter the name of the column which contains time zones for the content to be crawled.
changeDetectingColumns	Enter the names of the columns that Amazon Q will use to detect content changes. Amazon Q will re-index content when there is a change in any of these columns
allowedUsersColumns	Enter the name of the column which contains User IDs to be allowed access to content.
allowedGroupsColumn	Enter the name of the column which contains User IDs to be allowed access to content.

Configuration	Description
sourceURIColumn	Enter the name of the column which contains Source URLs to be indexed.
isSslEnabled	true to add a path to an SSL certificate file stored in an Amazon S3 bucket.
type	The type of data source. Specify JDBC as your data source type.
syncMode	<p>Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose</p> <ul style="list-style-type: none"> • <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index • <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index • <code>CHANGE_LOG</code> to incrementally crawl only new and modified content each time your data source syncs with your index.
secretArn	<p>The Amazon Resource Name (ARN) of a Secrets Manager secret that contains user name and password required to connect to your database. The secret must contain a JSON structure with the following keys:</p> <pre data-bbox="829 1629 1507 1829"> { "user name": "<i>database user name</i>", "password": "<i>password</i>" } </pre>

Configuration	Description
version	The version of the template that is currently supported.

How Amazon Q Business connector crawls Microsoft SQL Server ACLs

When you connect a database data source to Amazon Q, Amazon Q crawls user and group information from a column in the source table. You specify this column in the console or using the `configuration` parameter as part of the `CreateDataSource` operation.

If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

A database data source has the following limitations:

- You can only specify an allow list for a database data source. You can't specify a deny list.
- You can only specify groups. You can't specify individual users for the allow list.
- The database column should be a string containing a semicolon delimited list of groups.

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business Microsoft SQL Server data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.

- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a attribute mapping already available, or if you want to map additional document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q PostgreSQL connector supports the following field mappings:

Supported field mappings

- [Document](#)

Document

JDBC field name	Index field name	Description	Data type
jd_document_id	jd_document_id	Custom	String
jd_document_title	jd_document_title	Custom	String
jd_source_uri	_source_uri	Default	String

IAM role for Amazon Q Business Microsoft SQL Server connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the `BatchPutDocument` and `BatchDeleteDocument` operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- Permission to access the SSL certificate stored in your Amazon S3 bucket.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AllowsAmazonQToGetS3Objects",
    "Action": [
      "s3:GetObject"
    ],
    "Resource": [
      "arn:aws:s3:::{{input_bucket_name}}/*"
    ],
    "Effect": "Allow",
    "Condition": {
      "StringEquals": {
        "aws:ResourceAccount": "{{account_id}}"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToGetSecret",
```



```

    "Effect": "Allow",
    "Action": [
      "secretsmanager:GetSecretValue"
    ],
    "Resource": [
      "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[{{secret_id}}]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToDecryptSecret",
    "Effect": "Allow",
    "Action": [
      "kms:Decrypt"
    ],
    "Resource": [
      "arn:aws:kms:{{region}}:{{account_id}}:key/[{{key_id}}]"
    ],
    "Condition": {
      "StringLike": {
        "kms:ViaService": [
          "secretsmanager.*.amazonaws.com"
        ]
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToIngestDocuments",
    "Effect": "Allow",
    "Action": [
      "qbusiness:BatchPutDocument",
      "qbusiness:BatchDeleteDocument"
    ],
    "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
    {{application_id}}/index/{{index_id}}"
  },
  {
    "Sid": "AllowsAmazonQToIngestPrincipalMapping",
    "Effect": "Allow",
    "Action": [
      "qbusiness:PutGroup",
      "qbusiness:CreateUser",
      "qbusiness>DeleteGroup",
      "qbusiness:UpdateUser",
      "qbusiness:ListGroup"
    ]
  }
}

```

```

    ],
    "Resource": [
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}",
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}",
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}/data-source/*"
    ]
},
{
    "Sid": "AllowsAmazonQToCreateAndDeleteNI",
    "Effect": "Allow",
    "Action": [
        "ec2:CreateNetworkInterface",
        "ec2:DeleteNetworkInterface"
    ],
    "Resource": [
        "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
        "arn:aws:ec2:{{region}}:{{account_id}}:security-group/
[{{security_group}}]"
    ]
},
{
    "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
    "Effect": "Allow",
    "Action": [
        "ec2:CreateNetworkInterface",
        "ec2:DeleteNetworkInterface"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
        "StringLike": {
            "aws:RequestTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
        },
        "ForAllValues:StringEquals": {
            "aws:TagKeys": [
                "AMAZON_Q"
            ]
        }
    }
},
{

```

```

        "Sid": "AllowsAmazonQToCreateTags",
        "Effect": "Allow",
        "Action": [
            "ec2:CreateTags"
        ],
        "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
        "Condition": {
            "StringEquals": {
                "ec2:CreateAction": "CreateNetworkInterface"
            }
        }
    },
    {
        "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
        "Effect": "Allow",
        "Action": [
            "ec2:CreateNetworkInterfacePermission"
        ],
        "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
        "Condition": {
            "StringLike": {
                "aws:ResourceTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
            }
        }
    },
    {
        "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
        "Effect": "Allow",
        "Action": [
            "ec2:DescribeNetworkInterfaces",
            "ec2:DescribeAvailabilityZones",
            "ec2:DescribeNetworkInterfaceAttribute",
            "ec2:DescribeVpcs",
            "ec2:DescribeRegions",
            "ec2:DescribeNetworkInterfacePermissions",
            "ec2:DescribeSubnets"
        ],
        "Resource": "*"
    }
]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToAssumeRoleForServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "{{source_account}}"
        },
        "ArnLike": {
          "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
        }
      }
    }
  ]
}
```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Known limitations for the Amazon Q Business Microsoft SQL Server connector

- Deleted database rows will not be tracked in when Amazon Q checks for updated content.
- The size of field names and values in a row of your database can't exceed 400KB.
- If you have a large amount of data in your database data source, and do not want Amazon Q to index all your database content after the first sync, you can choose to sync only new, modified, or deleted documents.

Connecting Microsoft Teams to Amazon Q Business

Microsoft Teams is an enterprise collaboration tool for messaging, meetings, and file sharing. You can connect Microsoft Teams instance to Amazon Q Business—using either the AWS Management Console or the [CreateDataSource](#) API—and create an Amazon Q web experience..

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).
- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [Microsoft Teams connector overview](#)
- [Prerequisites for connecting Amazon Q Business to Microsoft Teams](#)
- [Connecting Amazon Q Business to Microsoft Teams using the console](#)
- [Connecting Amazon Q Business to Microsoft Teams using APIs](#)
- [How Amazon Q Business connector crawls Microsoft Teams ACLs](#)
- [Amazon Q Business Microsoft Teams data source connector field mappings](#)
- [IAM role for Amazon Q Business Microsoft Teams connector](#)
- [Troubleshooting your Amazon Q Business Microsoft Teams connector](#)

Microsoft Teams connector overview

The following table gives an overview of the Amazon Q Business Microsoft Teams connector and its supported features.

Category	Feature	Support
Security	Authentication type	OAuth 2.0
	Authentication credentials	<ul style="list-style-type: none"> • Microsoft Teams Client ID • Microsoft Teams Client secret

Category	Feature	Support
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Identity crawling	Yes
	VPC	Yes
Crawl features	Custom metadata	Yes
	Entities	<p>Yes. The following entities are supported:</p> <ul style="list-style-type: none"> • Chat messages • Chat attachments • Channel posts • Channel file attachments • Wiki • Meeting chats • Meeting details • Meeting notes • Meeting files
	Field mappings	Yes. Supports both default and custom field mappings. For more information, see Field mappings .

Category	Feature	Support
	Filters	<p>Yes. The following filters are supported:</p> <ul style="list-style-type: none"> • Include/exclude using user email • Include/exclude using team name • Include/exclude using channel name • Include/exclude using file name • Include/exclude using file type • Chat message • Chat attachment • Channel post • Channel attachment • Channel wiki • Calendar meeting • Meeting chat • Meeting file • Meeting note
	Sync mode	Supports full and incremental sync.
	File types	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to Microsoft Teams

Before you begin, make sure that you have completed the following prerequisites.

In Microsoft Teams, make sure you have:

- Created a Microsoft Teams account in Office 365.
- Copied your Microsoft 365 tenant ID. You can find your tenant ID in the Properties of your Azure Active Directory Portal. You need this URL to allow Amazon Q to connect with your Microsoft Teams data source.

- Configured an OAuth 2.0 credential token containing a client ID and client secret. For more information, see [Azure documentation on managing access tokens for Teams](#) on the Microsoft website.
- Added the necessary permissions. You can choose to add all permissions, or you can limit the scope by selecting fewer permissions based on which entities you want to crawl. The following table shows permissions by corresponding entity.

Entity	Required permissions for data sync	Required permissions for identity sync
Channel Post	<ul style="list-style-type: none"> • ChannelMessage.Read.All • Group.Read.All • User.Read • User.Read.All 	TeamMember.Read.All
Channel Attachment	<ul style="list-style-type: none"> • ChannelMessage.Read.All • Group.Read.All • User.Read • User.Read.All 	TeamMember.Read.All
Channel Wiki	<ul style="list-style-type: none"> • Group.Read.All • User.Read • User.Read.All 	TeamMember.Read.All
Chat Message	<ul style="list-style-type: none"> • Chat.Read.All • ChatMessage.Read.All • ChatMember.Read.All • User.Read • User.Read.All • Group.Read.All 	TeamMember.Read.All
Meeting Chat	<ul style="list-style-type: none"> • Chat.Read.All • ChatMessage.Read • ChatMember.Read.All 	TeamMember.Read.All

Entity	Required permissions for data sync	Required permissions for identity sync
	<ul style="list-style-type: none"> • User.Read • User.Read.All • Group.Read.All 	
Chat Attachment	<ul style="list-style-type: none"> • Chat.Read.All • ChatMessage.Read • ChatMember.Read.All • User.Read • User.Read.All • Group.Read.All 	TeamMember.Read.All
Meeting File	<ul style="list-style-type: none"> • Chat.Read.All • ChatMessage.Read.All • ChatMember.Read.All • User.Read • User.Read.All • Group.Read.All • Files.Read.All 	TeamMember.Read.All
Calendar Meeting	<ul style="list-style-type: none"> • Chat.Read.All • ChatMessage.Read.All • ChatMember.Read.All • User.Read • User.Read.All • Group.Read.All • Files.Read.All 	TeamMember.Read.All

Entity	Required permissions for data sync	Required permissions for identity sync
Meeting Notes	<ul style="list-style-type: none"> User.Read User.Read.All Group.Read.All Files.Read.All 	TeamMember.Read.All

- Generated Microsoft Teams OAuth 2.0 credentials containing a client id, client secret, username, and password. You need these credentials to authenticate Amazon Q to access Microsoft Teams.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your Microsoft Teams authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Connecting Amazon Q Business to Microsoft Teams using the console

The following procedure outlines how to connect Amazon Q Business to Microsoft Teams using the AWS Management Console.

Connecting Amazon Q to Microsoft Teams

- Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
- Complete the steps to [create your Amazon Q application](#).

3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **Microsoft Teams** page, enter the following information:
6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. In **Source**, enter the following information:
 - **Tenant ID** – Enter your tenant id. Your Microsoft tenant ID is a globally unique identifier that's necessary to configure each connector instance. Your tenant ID is different from your organization name or domain and can be found in the properties section of your Microsoft account dashboard.
8. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

 **Note**

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

9. **Authentication** – Choose between **New** and **Existing**.
 - If you choose **Existing**, select an existing secret for **Select secret**.

If you choose **New**, enter the following information in the **New AWS Secrets Manager secret** section:
 - i. **Secret name** – A name for your secret.
 - ii. For **Client ID, Client secret** – Enter the authentication credential values that you generated from your Teams account.
10. **Payment model** – You can choose a licensing and payment model for your Teams account. Model A payment models are restricted to licensing and payment models that require security

compliance. Model B payment models are suitable for licensing and payment models that don't require security compliance.

11. **Configure VPC and security group – optional** – Choose whether you want to use a VPC. If you do, enter the following information:
 - a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
 - b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

12. **Identity crawler** – Choose to activate Amazon Q identity crawler to sync identity information. For more information, see [Identity crawler](#).
13. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

14. **Sync scope** – Select the content you want to sync.
15. In **Additional configuration – optional**, choose from the following options:
 - **Calendar crawling** – Enter the date range for which the connector will crawl your calendar content.
 - **User email** – Enter the user emails you wish to include in your application.
 - **Team names** – Add patterns to include or exclude teams found in Microsoft Teams from your application.
 - **Channel names** – Add patterns to include or exclude channels found in Microsoft Teams from your application.
 - **Attachment regex patterns** – Add regular expression patterns to include or exclude certain attachment for all supported entities. You can add up to 100 patterns.
16. In **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.
 - **Full sync** – Sync all content regardless of the previous sync status.

- **New or modified content sync** – Sync only new and modified documents.
- **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.

For more details, see [Sync mode](#).

17. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
18. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
19. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:
 - a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
 - b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

 **Note**

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

20. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

21. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

 **Note**

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to

view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to Microsoft Teams using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the configuration parameter to provide a JSON schema with all other configuration information specific to your data source connector.

For an example of the API request, see [CreateDataSource](#) in the Amazon Q API Reference.

Microsoft Teams JSON schema

The following is the Microsoft Teams JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
      "type": "object",
      "properties": {
        "repositoryEndpointMetadata": {
          "type": "object",
          "properties": {
            "tenantId": {
              "type": "string",
              "pattern": "^[0-9a-f]{8}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{4}-[0-9a-f]{12}$",
              "minLength": 36,
              "maxLength": 36
            }
          },
          "required": [
            "tenantId"
          ]
        }
      },
      "required": [
        "repositoryEndpointMetadata"
      ]
    }
  }
}
```

```

},
"repositoryConfigurations": {
  "type": "object",
  "properties": {
    "chatMessage": {
      "type": "object",
      "properties": {
        "fieldMappings": {
          "type": "array",
          "items": [
            {
              "type": "object",
              "properties": {
                "indexFieldName": {
                  "type": "string"
                },
                "indexFieldType": {
                  "type": "string",
                  "enum": [
                    "STRING",
                    "STRING_LIST",
                    "DATE"
                  ]
                },
                "dataSourceFieldName": {
                  "type": "string"
                },
                "dateFieldFormat": {
                  "type": "string",
                  "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
                }
              }
            },
            "required": [
              "indexFieldName",
              "indexFieldType",
              "dataSourceFieldName"
            ]
          ]
        }
      }
    },
    "required": [
      "fieldMappings"
    ]
  }
}

```

```

    },
    "chatAttachment": {
      "type": "object",
      "properties": {
        "fieldMappings": {
          "type": "array",
          "items": [
            {
              "type": "object",
              "properties": {
                "indexFieldName": {
                  "type": "string"
                },
                "indexFieldType": {
                  "type": "string",
                  "enum": [
                    "STRING",
                    "DATE",
                    "LONG"
                  ]
                },
                "dataSourceFieldName": {
                  "type": "string"
                },
                "dateFieldFormat": {
                  "type": "string",
                  "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
                }
              }
            }
          ]
        },
        "required": [
          "indexFieldName",
          "indexFieldType",
          "dataSourceFieldName"
        ]
      }
    },
    "required": [
      "fieldMappings"
    ]
  },
  "channelPost": {
    "type": "object",

```



```

    "properties": {
      "fieldMappings": {
        "type": "array",
        "items": [
          {
            "type": "object",
            "properties": {
              "indexFieldName": {
                "type": "string"
              },
              "indexFieldType": {
                "type": "string",
                "enum": [
                  "STRING",
                  "STRING_LIST",
                  "DATE"
                ]
              },
              "dataSourceFieldName": {
                "type": "string"
              },
              "dateFieldFormat": {
                "type": "string",
                "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
              }
            },
            "required": [
              "indexFieldName",
              "indexFieldType",
              "dataSourceFieldName"
            ]
          }
        ]
      },
      "required": [
        "fieldMappings"
      ]
    },
    "channelWiki": {
      "type": "object",
      "properties": {
        "fieldMappings": {
          "type": "array",

```

```

    "items": [
      {
        "type": "object",
        "properties": {
          "indexFieldName": {
            "type": "string"
          },
          "indexFieldType": {
            "type": "string",
            "enum": [
              "STRING",
              "DATE",
              "LONG"
            ]
          },
          "dataSourceFieldName": {
            "type": "string"
          },
          "dateFieldFormat": {
            "type": "string",
            "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
          }
        },
        "required": [
          "indexFieldName",
          "indexFieldType",
          "dataSourceFieldName"
        ]
      }
    ],
    "required": [
      "fieldMappings"
    ],
    "channelAttachment": {
      "type": "object",
      "properties": {
        "fieldMappings": {
          "type": "array",
          "items": [
            {
              "type": "object",

```

```

        "properties": {
          "indexFieldName": {
            "type": "string"
          },
          "indexFieldType": {
            "type": "string",
            "enum": [
              "STRING",
              "DATE",
              "LONG"
            ]
          },
          "dataSourceFieldName": {
            "type": "string"
          },
          "dateFieldFormat": {
            "type": "string",
            "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
          }
        },
        "required": [
          "indexFieldName",
          "indexFieldType",
          "dataSourceFieldName"
        ]
      }
    ]
  },
  "required": [
    "fieldMappings"
  ]
},
"meetingChat": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": [
        {
          "type": "object",
          "properties": {
            "indexFieldName": {
              "type": "string"
            }
          }
        }
      ]
    }
  }
}

```

```

    },
    "indexFieldType": {
      "type": "string",
      "enum": [
        "STRING",
        "STRING_LIST",
        "DATE"
      ]
    },
    "dataSourceFieldName": {
      "type": "string"
    },
    "dateFieldFormat": {
      "type": "string",
      "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
    }
  ],
  "required": [
    "indexFieldName",
    "indexFieldType",
    "dataSourceFieldName"
  ]
}
]
}
},
"required": [
  "fieldMappings"
]
},
"meetingFile": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": [
        {
          "type": "object",
          "properties": {
            "indexFieldName": {
              "type": "string"
            },
            "indexFieldType": {
              "type": "string",

```

```

        "enum": [
            "STRING",
            "DATE",
            "LONG"
        ]
    },
    "dataSourceFieldName": {
        "type": "string"
    },
    "dateFieldFormat": {
        "type": "string",
        "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
    }
},
"required": [
    "indexFieldName",
    "indexFieldType",
    "dataSourceFieldName"
]
}
]
}
},
"required": [
    "fieldMappings"
]
},
"meetingNote": {
    "type": "object",
    "properties": {
        "fieldMappings": {
            "type": "array",
            "items": [
                {
                    "type": "object",
                    "properties": {
                        "indexFieldName": {
                            "type": "string"
                        },
                        "indexFieldType": {
                            "type": "string",
                            "enum": [
                                "STRING",
                                "DATE"
                            ]
                        }
                    }
                }
            ]
        }
    }
}

```

```

    ]
    },
    "dataSourceFieldName": {
      "type": "string"
    },
    "dateFieldFormat": {
      "type": "string",
      "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
    }
  ],
  "required": [
    "indexFieldName",
    "indexFieldType",
    "dataSourceFieldName"
  ]
}
]
}
},
"required": [
  "fieldMappings"
]
},
"calendarMeeting": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": [
        {
          "type": "object",
          "properties": {
            "indexFieldName": {
              "type": "string"
            },
            "indexFieldType": {
              "type": "string",
              "enum": [
                "STRING",
                "DATE"
              ]
            }
          }
        }
      ]
    },
    "dataSourceFieldName": {
      "type": "string"
    }
  }
}

```

```

        },
        "dateFieldFormat": {
            "type": "string",
            "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
        }
    },
    "required": [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
    ]
}
]
}
},
"required": [
    "fieldMappings"
]
}
},

```

```

"additionalProperties": {
    "type": "object",
    "properties": {
        "isCrawlAcl": {
            "type": "boolean"
        },
        "fieldForUserId": {
            "type": "string"
        },
        "paymentModel": {
            "type": "string",
            "enum": [
                "A",
                "B",
                "Evaluation Mode"
            ]
        },
        "inclusionTeamNameFilter": {
            "type": "array",
            "items": {
                "type": "string"
            }
        }
    }
}

```

```
    }
  },
  "exclusionTeamNameFilter": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "inclusionChannelNameFilter": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "exclusionChannelNameFilter": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "inclusionFileNamePatterns": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "exclusionFileNamePatterns": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "inclusionFileTypePatterns": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "exclusionFileTypePatterns": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
},
```



```
"inclusionUserEmailFilter": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"isCrawlChatMessage": {
  "type": "boolean"
},
"isCrawlChatAttachment": {
  "type": "boolean"
},
"isCrawlChannelPost": {
  "type": "boolean"
},
"isCrawlChannelAttachment": {
  "type": "boolean"
},
"isCrawlChannelWiki": {
  "type": "boolean"
},
"isCrawlCalendarMeeting": {
  "type": "boolean"
},
"isCrawlMeetingChat": {
  "type": "boolean"
},
"isCrawlMeetingFile": {
  "type": "boolean"
},
"isCrawlMeetingNote": {
  "type": "boolean"
},
"startCalendarDateTime": {
  "anyOf": [
    {
      "type": "string",
      "pattern": "^[0-9]{4}-[0-9]{2}-[0-9]{2}T[0-9]{2}:[0-9]{2}:[0-9]{2}Z$"
    },
    {
      "type": "string",
      "pattern": ""
    }
  ]
}
```

```

    },
    "endCalendarDateTime": {
      "anyOf": [
        {
          "type": "string",
          "pattern": "^[0-9]{4}-[0-9]{2}-[0-9]{2}T[0-9]{2}:[0-9]{2}:[0-9]{2}Z$"
        },
        {
          "type": "string",
          "pattern": ""
        }
      ]
    },
    "required": []
  },
  "type": {
    "type": "string",
    "pattern": "MSTEAMS"
  },
  "enableIdentityCrawler": {
    "type": "boolean"
  },
  "syncMode": {
    "type": "string",
    "enum": [
      "FORCED_FULL_CRAWL",
      "FULL_CRAWL",
      "CHANGE_LOG"
    ]
  },
  "secretArn": {
    "type": "string",
    "minLength": 20,
    "maxLength": 2048
  }
},
"version": {
  "type": "string",
  "anyOf": [
    {
      "pattern": "1.0.0"
    }
  ]
}
]

```

```

},
"required": [
  "connectionConfiguration",
  "repositoryConfigurations",
  "syncMode",
  "additionalProperties",
  "secretArn",
  "type"
]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for endpoint for the data source.
repositoryEndpointMetadata	The endpoint information for the data source.
tenantId	The Microsoft 365 tenant ID. You can find your tenant ID in the Properties of your Azure Active Directory Portal.
repositoryConfigurations	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings.
<ul style="list-style-type: none"> • chatMessage • chatAttachment • channelPost • channelWiki • channelAttachment • meetingChat • meetingFile • meetingNote • calendarMeeting 	A list of objects that map the attributes or field names of your Microsoft Teams content to Amazon Q index field names.

Configuration	Description
additionalProperties	Additional configuration options for your content in your data source.
isCrawlAcl	Specify true to crawl access control information from documents.
fieldForUserId	Specify field to use for UserId for ACL crawling.
<ul style="list-style-type: none"> • isCrawlChatMessage • isCrawlChatAttachment • isCrawlChannelPost • isCrawlChannelAttachment • isCrawlChannelWiki • isCrawlCalendarMeeting • isCrawlMeetingChat • isCrawlMeetingFile • isCrawlMeetingNote 	true to index this content in your Microsoft Teams data source.
paymentModel	Specifies what type of payment model to use with your Teams data source. Model A payment models are restricted to licensing and payment models that require security compliance. Model B payment models are suitable for licensing and payment models that don't require security compliance.

Configuration	Description
syncMode	<p>Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose between the following options:</p> <ul style="list-style-type: none">• Use <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index• Use <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index• Use <code>CHANGE_LOG</code> to incrementally crawl only new and modified content each time your data source syncs with your index
secretARN	<p>The Amazon Resource Name (ARN) of an AWS Secrets Manager secret that contains the key-value pairs required to connect to your Microsoft Teams. This includes your client ID and client secret.</p>
type	<p>The type of data source. Specify <code>MSTEAMS</code> as your data source type.</p>
enableIdentityCrawler	<p><code>true</code> to activate identity crawler. Crawling identity information on users and groups with access to specific documents is useful for user context filtering. Search results are filtered based on the user or their group access to documents. See Identity crawler for more information.</p>

Configuration	Description
version	The version of this template that's currently supported.

How Amazon Q Business connector crawls Microsoft Teams ACLs

When you connect an Microsoft Teams data source to Amazon Q Business, Amazon Q Business crawls ACL information attached to a document (user and group information) from your Microsoft Teams instance. If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

The group and user IDs are mapped as follows:

- `_tenant_id` – Your Microsoft tenant ID is a globally unique identifier that's necessary to configure each connector instance. Your tenant ID is different from your organization name or domain and can be found in the properties section of your Microsoft account dashboard.
- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business Microsoft Teams data source connector field mappings

To help you structure data for retrieval and chat filtering, Amazon Q Business crawls data source document attributes or metadata and maps them to fields in your Amazon Q index.

Amazon Q has reserved fields that it uses when querying your application. When possible, Amazon Q automatically maps these built-in fields to attributes in your data source. If a built-in field doesn't have a default mapping, or if you want to map additional index fields, use the custom field mappings to specify how a data source attribute maps to your Amazon Q application. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

⚠ Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q Teams connector supports the following entities and the associated reserved and custom attributes.

ℹ Note

You can map any Teams field to the document title or document body Amazon Q reserved/default index fields.

Supported entities and field mappings

- [Chat messages](#)
- [Chat attachments](#)
- [Channel posts](#)
- [Channel file attachments](#)
- [Wiki](#)
- [Meeting chats](#)
- [Meeting details](#)
- [Meeting notes](#)
- [Meeting files](#)

Chat messages

Microsoft Teams field name	Index field name	Description	Data type
subject	tms_subject	Custom	String
summary	tms_summary	Custom	String

Microsoft Teams field name	Index field name	Description	Data type
importance	tms_importance	Custom	String
messageType	tms_message_type	Custom	String
sender	tms_sender	Custom	String
sourceUrl	_source_uri	Default	String
attachments	tms_attachments	Custom	String list
reactions	tms_reactions	Custom	String list
mentions	tms_mentions	Custom	String list
deletedAt	tms_last_deleted_at	Custom	Date
createdAt	_created_at	Default	Date
lastModifiedAt	_last_updated_at	Default	Date

Chat attachments

Microsoft Teams field name	Index field name	Description	Data type
fileName	tms_name	Custom	String
size	tms_file_size	Custom	Long (numeric)
title	tms_title	Custom	String
sourceUrl	_source_uri	Default	String
lastModifiedBy	tms_last_modified_by	Custom	String
createdBy	tms_created_by	Custom	String

Microsoft Teams field name	Index field name	Description	Data type
createdAt	_created_at	Default	Date
lastModifiedAt	_last_updated_at	Default	Date

Channel posts

Microsoft Teams field name	Index field name	Description	Data type
subject	tms_subject	Custom	String
summary	tms_summary	Custom	String
importance	tms_importance	Custom	String
messageType	tms_message_type	Custom	String
createdBy	tms_created_by	Custom	String
deletedAt	tms_last_deleted_at	Custom	Date
sourceUrl	_source_uri	Default	String
mentions	tms_mentions	Custom	String list
reactions	tms_reactions	Custom	String list
attachments	tms_attachments	Custom	String list
createdAt	_created_at	Default	Date
lastModifiedAt	_last_updated_at	Default	Date

Channel file attachments

Microsoft Teams field name	Index field name	Description	Data type
fileName	tms_name	Custom	String
size	tms_file_size	Custom	Long (numeric)
channelName	tms_channel_name	Custom	String
Title	tms_title	Custom	String
sourceUrl	_source_uri	Default	String
createdBy	tms_created_by	Custom	String
lastModifiedBy	tms_last_modified_by	Custom	String
createdAt	_created_at	Default	Date
lastModifiedAt	_last_updated_at	Default	Date
oneNoteDocument	tms_onenote_document	Custom	String
oneNoteSection	tms_onenote_section	Custom	String
oneNotePage	tms_onenote_page	Custom	String

Wiki

Microsoft Teams field name	Index field name	Description	Data type
channelName	tms_channel_name	Custom	String
fileName	tms_name	Custom	String

Microsoft Teams field name	Index field name	Description	Data type
size	tms_file_size	Custom	Long (numeric)
createdBy	tms_created_by	Custom	String
lastModifiedBy	tms_last_modified_by	Custom	String
title	tms_title	Custom	String
sourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
lastModifiedAt	_last_updated_at	Default	Date

Meeting chats

Microsoft Teams field name	Index field name	Description	Data type
subject	tms_subject	Custom	String
summary	tms_summary	Custom	String
importance	tms_importance	Custom	String
messageType	tms_message_type	Custom	String
Sender	tms_sender	Custom	String
attachments	tms_attachments	Custom	String list
mentions	tms_mentions	Custom	String list
reactions	tms_reactions	Custom	String list
sourceUrl	_source_uri	Default	String

Microsoft Teams field name	Index field name	Description	Data type
deletedAt	tms_last_deleted_at	Custom	Date
createdAt	_created_at	Default	Date
lastModifiedAt	_last_updated_at	Default	Date

Meeting details

Microsoft Teams field name	Index field name	Description	Data type
subject	tms_subject	Custom	String
summary	tms_summary	Custom	String
importance	tms_importance	Custom	String
username	tms_from_user	Custom	String
eventStartTime	tms_event_start_time	Custom	Date
eventEndTime	tms_event_end_time	Custom	Date
sourceURL	_source_uri	Default	String

Meeting notes

Microsoft Teams field name	Index field name	Description	Data type
fileName	tms_name	Custom	String
title	tms_title	Custom	String
createdBy	tms_created_by	Custom	String

Microsoft Teams field name	Index field name	Description	Data type
lastModifiedBy	tms_last_modified_by	Custom	String
sourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
lastModifiedAt	_last_updated_at	Default	Date

Meeting files

Microsoft Teams field name	Index field name	Description	Data type
fileName	tms_name	Custom	String
title	tms_title	Custom	String
size	tms_file_size	Custom	Long (numeric)
sourceUrl	_source_uri	Default	String
createdBy	tms_created_by	Custom	String
lastModifiedBy	tms_last_modified_by	Custom	String
createdAt	_created_at	Default	Date
lastModifiedAt	_last_updated_at	Default	Date

IAM role for Amazon Q Business Microsoft Teams connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the `BatchPutDocument` and `BatchDeleteDocument` operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToGetSecret",
      "Effect": "Allow",
      "Action": [
        "secretsmanager:GetSecretValue"
      ],
      "Resource": [
        "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
      ]
    },
    {
      "Sid": "AllowsAmazonQToDecryptSecret",
      "Effect": "Allow",
      "Action": [
        "kms:Decrypt"
      ],
      "Resource": [
        "arn:aws:kms:{{region}}:{{account_id}}:key/[key_id]"
      ],
      "Condition": {
        "StringLike": {
          "kms:ViaService": [
            "secretsmanager.*.amazonaws.com"
          ]
        }
      }
    }
  ]
}
```

```

    ]
  }
}
},
{
  "Sid": "AllowsAmazonQToIngestDocuments",
  "Effect": "Allow",
  "Action": [
    "qbusiness:BatchPutDocument",
    "qbusiness:BatchDeleteDocument"
  ],
  "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
},
{
  "Sid": "AllowsAmazonQToIngestPrincipalMapping",
  "Effect": "Allow",
  "Action": [
    "qbusiness:PutGroup",
    "qbusiness:CreateUser",
    "qbusiness>DeleteGroup",
    "qbusiness:UpdateUser",
    "qbusiness:ListGroups"
  ],
  "Resource": [
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}",
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}",
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}/data-source/*"
  ]
},
{
  "Sid": "AllowsAmazonQToCreateAndDeleteNI",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateNetworkInterface",
    "ec2>DeleteNetworkInterface"
  ],
  "Resource": [
    "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
    "arn:aws:ec2:{{region}}:{{account_id}}:security-group/[{{security_group}}]"
  ]
},

```

```

{
  "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateNetworkInterface",
    "ec2:DeleteNetworkInterface"
  ],
  "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
  "Condition": {
    "StringLike": {
      "aws:RequestTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
    },
    "ForAllValues:StringEquals": {
      "aws:TagKeys": [
        "AMAZON_Q"
      ]
    }
  }
},
{
  "Sid": "AllowsAmazonQToCreateTags",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateTags"
  ],
  "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
  "Condition": {
    "StringEquals": {
      "ec2:CreateAction": "CreateNetworkInterface"
    }
  }
},
{
  "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateNetworkInterfacePermission"
  ],
  "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
  "Condition": {
    "StringLike": {
      "aws:ResourceTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
    }
  }
}

```



```

    },
    {
      "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
      "Effect": "Allow",
      "Action": [
        "ec2:DescribeNetworkInterfaces",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeNetworkInterfaceAttribute",
        "ec2:DescribeVpcs",
        "ec2:DescribeRegions",
        "ec2:DescribeNetworkInterfacePermissions",
        "ec2:DescribeSubnets"
      ],
      "Resource": "*"
    }
  ]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "{{source_account}}"
        },
        "ArnEquals": {
          "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
        }
      }
    }
  ]
}

```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Troubleshooting your Amazon Q Business Microsoft Teams connector

The following table provides information about error codes you may see for the Microsoft Teams connector and suggested troubleshooting actions.

Error code	Error message	Suggested resolution
MST-5001	Exception occurred while sending request to MSTeams api, please try again later.	Error related to authentication. Check logs for the specific error message.
MST-5101	Exception occurred while validating configuration.	Error related to configurations. Check logs for the specific error message.
MST-5102	ClientID cannot be null in Repository configuration.	Error related to configurations. Check logs for the specific error message.
MST-5103	TenantId cannot be null in Repository configuration.	Error related to configurations. Check logs for the specific error message.
MST-5104	ClientSecret cannot be null in Repository configuration	Error related to configurations. Check logs for the specific error message.
MST-5105	Please add a valid paymentModel under additionalProperties. The paymentModel should be one of the following.	Error related to configurations. Check logs for the specific error message.
MST-5106	Please add valid startCalendarDateTime & endCalendarDateTime under additionalProperti	Error related to configurations. Please check logs for the specific error message.

Error code	Error message	Suggested resolution
	es: startCalendarDateT ime & endCalend arDateTime should be in this format 2016-12-0 1T00:00:00Z.	
MST-5107	isCrawlChatMessage should be true or false.	Error related to configurations. Please check logs for the specific error message.
MST-5108	isCrawlMeetingChatValue should be true or false.	Error related to configurations. Please check logs for the specific error message.
MST-5109	isCrawlChatAttachment should be true or false.	Error related to configurations. Please check logs for the specific error message.
MST-5110	isCrawlMeetingFile should be true or false.	Error related to configurations. Please check logs for the specific error message.
MST-5111	isCrawlMeetingNote should be true or false.	Error related to configurations. Please check logs for the specific error message.
MST-5112	isCrawlChannelPost should be true or false.	Error related to configurations. Please check logs for the specific error message.
MST-5113	isCrawlChannelAtta chment should be true or false	Error related to configurations. Please check logs for the specific error message.
MST-5114	isCrawlChannelWiki should be true or false.	Error related to configurations. Please check logs for the specific error message.

Error code	Error message	Suggested resolution
MST-5115	isCrawlCalendarMeeting should be true or false.	Error related to configurations. Please check logs for the specific error message.
MST-5116	Invalid clientId pattern.	Error related to configurations. Please check logs for the specific error message..
MST-5117	ClientSecret Over maximum length.	Error related to configurations. Please check logs for the specific error message.
MST-5200	Got exception from customer while accessing list of users.	Failure while fetching the list of users from Microsoft Graph API. Please check logs for more details.
MST-5201	Got exception from customer while accessing list of chats.	Failure while fetching the list of chats from Microsoft Graph API. Please check logs for more details.
MST-5202	Got exception from customer while accessing meeting chats.	Failure while fetching meeting chats from Microsoft Graph API. Please check logs for more details.
MST-5203	Got exception from customer while accessing list of groups.	Failure while fetching the list of groups from Microsoft Graph API. Please check logs for more details.
MST-5204	Got exception from customer while accessing list of channels.	Failure while fetching the list of channels from Microsoft Graph API. Please check logs for more details.
MST-5205	Error occurred while fetching meeting events.	Failure while fetching meeting events from Microsoft Graph API. Please check logs for more details.

Error code	Error message	Suggested resolution
MST-5206	Error occurred while fetching drive files.	Failure while fetching drive files from Microsoft Graph API. Please check logs for more details.
MST-5207	Error while InterruptedException rate limit.	Failures while retrying API requests to fetch data from Microsoft Graph API.
MST-5209	Got exception from customer while running full crawl.	Failures while running full crawl iterator. Please refer logs or contact connector team for more information.
MST-5210	Exception occurred while accessing list of channel attachment from data source.	Failure while fetching the list of channels attachment from Microsoft Graph API. Please check logs for more details.
MST-5211	Exception occurred while accessing meeting chat information for user.	Failure while accessing meeting chats from Microsoft Graph API. Please check logs for more details.
MST-5212	Exception occurred while processing to access list of users.	Failure while processing to access list of users from Microsoft Graph API. Please check logs for more details.
MST-5213	Exception occurred while processing to access list of groups.	Failure while processing to access list of groups from Microsoft Graph API. Please check logs for more details.
MST-5214	Exception occurred while processing to access list of channel attachment.	Failure while processing to access list of channel attachment from Microsoft Graph API. Please check logs for more details.
MST-5215	Exception occurred while processing to access meeting events.	Failure while processing to access meeting events from Microsoft Graph API. Please check logs for more details.

Error code	Error message	Suggested resolution
MST-5301	Got exception from customer while running changelog.	Failures while handling changelog token. Please refer logs or contact connector team for more information.
MST-5302	Error in serializing change log token.	Failures while serializing change log token. Please refer logs or contact connector team for more information.
MST-5303	Error in de-serializing change log token.	Failures while de-serializing change log token. Please refer logs or contact connector team for more information.
MST-5400	Exception occurred while running Identity Crawler.	Error occurred while fetching groups details from Microsoft Graph API. Please check logs for more details.
MST-5401	Error while build groups details for Identity Crawler.	Failures while de-serializing change log token. Please refer logs or contact connector team for more information.
MST-5500	Exception occurred while getting file content response.	Error occurred while fetching file content response details from Microsoft Graph API. Please check logs for more details.
MST-5501	Only String, String List, Date and Long formats are supported for field mappings.	Error related to unsupported field mappings. Please check logs for the specific error message.
MST-5502	IO Exception occurred.	IO Exception.

Connecting Microsoft Yammer to Amazon Q Business

Microsoft Yammer is an enterprise collaboration tool for messaging, meetings, and file sharing. You can connect Microsoft Yammer instance to Amazon Q Business—using either the AWS Management Console or the [CreateDataSource](#) API—and create an Amazon Q web experience.

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).
- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [Microsoft Yammer connector overview](#)
- [Prerequisites for connecting Amazon Q to Microsoft Yammer](#)
- [Connecting Amazon Q Business to Microsoft Yammer using the console](#)
- [Connecting Amazon Q Business to Microsoft Yammer using APIs](#)
- [How Amazon Q Business connector crawls Microsoft Yammer ACLs](#)
- [Amazon Q Business Microsoft Yammer data source connector field mappings](#)
- [IAM role for Amazon Q Business Microsoft Yammer connector](#)
- [Known limitations for the Amazon Q Business Microsoft Yammer connector](#)
- [Troubleshooting your Amazon Q Business Microsoft Yammer connector](#)

Microsoft Yammer connector overview

The following table gives an overview of the Amazon Q Business Microsoft Yammer connector and its supported features.

Category	Feature	Support
Security	Authentication type	OAuth 2.0
	Authentication credentials	<ul style="list-style-type: none"> • Microsoft Yammer username

Category	Feature	Support
		<ul style="list-style-type: none"> • Microsoft Yammer password • Microsoft Yammer Client ID • Microsoft Yammer Client secret
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Identity crawling	Yes
	VPC	Yes
Crawl features	Custom metadata	Yes
	Entities	Yes. The following entities are supported: <ul style="list-style-type: none"> • Message • Attachment • User • Community
	Field mappings	Yes. Supports both default and custom field mappings. For more information, see Field mappings .
	Filters	Yes. The following filters are supported: <ul style="list-style-type: none"> • Community names • Public messages • Attachments • Inbox private messages • Crawl content beginning from a date • Including and excluding content by file type
	Sync mode	Supports full and incremental sync.
	File types	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q to Microsoft Yammer

Before you begin, make sure that you have completed the following prerequisites.

In Microsoft Yammer, make sure you have:

- Created a Microsoft Yammer administrative account with verified admin user permissions.
- Configured an OAuth 2.0 credential token containing a client ID and client secret.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your Microsoft Yammer authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Connecting Amazon Q Business to Microsoft Yammer using the console

The following procedure outlines how to connect Amazon Q Business to Microsoft Yammer using the AWS Management Console.

Connecting Amazon Q to Microsoft Yammer

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **Microsoft Yammer** page, enter the following information:
6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

 **Note**

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

8. **Authentication** – Choose between **New** and **Existing**.

- If you choose **Existing**, select an existing secret for **Select secret**.

If you choose **New**, enter the following information in the **New AWS Secrets Manager secret** section:

- **Secret name** – A name for your secret.
- **Username** – The username for your Microsoft Yammer Active Directory account.
- **Password** – The password for your Microsoft Yammer Active Directory account.
- **Client ID** – The OAuth client ID credential values you copied from your Microsoft Yammer account.
- **Client secret** – The client secret from your Microsoft Yammer account.

9. **Configure VPC and security group** – *optional* – Choose whether you want to use a VPC. If you do, enter the following information:

- a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
- b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

10. **Identity crawler** – Choose to activate Amazon Q identity crawler to sync identity information. For more information, see [Identity crawler](#).

11. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

12. For **Sync scope**, provide the following information:

- **sinceDate** – Select the date in your data source content from when Amazon Q should begin to crawl your data.
- **Select content to sync** – Choose between **All**, **Public messages**, **Attachments**, and **Inbox private messages**.

13. For **Additional configuration – optional**, provide the following information:

- **Community names** – Enter the community names you wish to include in your application.
- **Regex patterns** – Add regular expression patterns to include or exclude certain file types. You can add up to 100 patterns.

14. In **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.

- **Full sync** – Sync all content regardless of the previous sync status.
- **New or modified content sync** – Sync only new and modified documents.
- **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.

For more details, see [Sync mode](#).

15. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).

16. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.

17. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:

- a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
- b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

 **Note**

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

18. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

19. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

 **Note**

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to Microsoft Yammer using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the configuration parameter to provide a JSON schema with all other configuration information specific to your data source connector.

For an example of the API request, see [CreateDataSource](#) in the Amazon Q API Reference.

Yammer JSON schema

The following is the Yammer JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
      "type": "object",
      "properties": {
        "repositoryEndpointMetadata": {
          "type": "object",
          "properties": {
          }
        }
      }
    },
    "required": [
      "repositoryEndpointMetadata"
    ]
  },
  "repositoryConfigurations": {
    "type": "object",
    "properties": {
      "community": {
        "type": "object",
        "properties": {
          "fieldMappings": {
            "type": "array",
            "items": {
              "anyOf": [
                {
                  "type": "object",
                  "properties": {
                    "indexFieldName": {
                      "type": "string"
                    },
                    "indexFieldType": {
                      "type": "string",
                      "enum": [
                        "STRING",
                        "DATE"
                      ]
                    }
                  }
                }
              ]
            }
          }
        }
      }
    }
  }
}
```

```

        "dataSourceFieldName": {
            "type": "string"
        },
        "dateFieldFormat": {
            "type": "string",
            "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
        }
    },
    "required": [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
    ]
}
]
}
},
"required": [
    "fieldMappings"
]
},
"user": {
    "type": "object",
    "properties": {
        "fieldMappings": {
            "type": "array",
            "items": {
                "anyOf": [
                    {
                        "type": "object",
                        "properties": {
                            "indexFieldName": {
                                "type": "string"
                            },
                            "indexFieldType": {
                                "type": "string",
                                "enum": [
                                    "STRING",
                                    "DATE"
                                ]
                            }
                        }
                    }
                ]
            }
        },
        "dataSourceFieldName": {
            "type": "string"
        }
    }
}

```

```

        },
        "dateFieldFormat": {
            "type": "string",
            "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
        }
    },
    "required": [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
    ]
}
]
}
},
"required": [
    "fieldMappings"
]
},
"message": {
    "type": "object",
    "properties": {
        "fieldMappings": {
            "type": "array",
            "items": {
                "anyOf": [
                    {
                        "type": "object",
                        "properties": {
                            "indexFieldName": {
                                "type": "string"
                            },
                        },
                        "indexFieldType": {
                            "type": "string",
                            "enum": [
                                "STRING",
                                "DATE"
                            ]
                        },
                    },
                    "dataSourceFieldName": {
                        "type": "string"
                    },
                ],
                "dateFieldFormat": {

```

```

        "type": "string",
        "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
    }
},
"required": [
    "indexFieldName",
    "indexFieldType",
    "dataSourceFieldName"
]
}
]
}
},
"required": [
    "fieldMappings"
]
},
"attachment": {
    "type": "object",
    "properties": {
        "fieldMappings": {
            "type": "array",
            "items": {
                "anyOf": [
                    {
                        "type": "object",
                        "properties": {
                            "indexFieldName": {
                                "type": "string"
                            },
                            "indexFieldType": {
                                "type": "string",
                                "enum": [
                                    "STRING",
                                    "DATE"
                                ]
                            }
                        },
                        "dataSourceFieldName": {
                            "type": "string"
                        },
                        "dateFieldFormat": {
                            "type": "string",
                            "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
                        }
                    }
                ]
            }
        }
    }
}

```



```

        }
      },
      "required": [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
      ]
    }
  ]
}
},
"required": [
  "fieldMappings"
]
}
},
"additionalProperties": {
  "type": "object",
  "properties": {
    "isCrawlAcl": {
      "type": "boolean"
    },
    "fieldForUserId": {
      "type": "string"
    },
    "inclusionPatterns": {
      "type": "array"
    },
    "exclusionPatterns": {
      "type": "array"
    },
    "sinceDate": {
      "type": "string",
      "pattern": "^(19|2[0-9])[0-9]{2}-(0[1-9]|1[012])-(0[1-9]|[12][0-9]|3[01])T(0[0-9]|1[0-9]|2[0-3]):([0-5][0-9]):([0-5][0-9])(\\+|-)(0[0-9]|1[0-9]|2[0-3]):([0-5][0-9]))?$"
    },
    "communityNameFilter": {
      "type": "array",
      "items": {
        "type": "string"
      }
    }
  }
}

```

```
    },
    "isCrawlMessage": {
      "type": "boolean"
    },
    "isCrawlAttachment": {
      "type": "boolean"
    },
    "isCrawlPrivateMessage": {
      "type": "boolean"
    }
  },
  "required": [
    "sinceDate"
  ]
},
"type": {
  "type": "string",
  "pattern": "YAMMER"
},
"secretArn": {
  "type": "string",
  "minLength": 20,
  "maxLength": 2048
},
"useChangeLog": {
  "type": "string",
  "enum": [
    "true",
    "false"
  ]
},
"syncMode": {
  "type": "string",
  "enum": [
    "FORCED_FULL_CRAWL",
    "FULL_CRAWL",
    "CHANGE_LOG"
  ]
},
"enableIdentityCrawler": {
  "type": "boolean"
},
"version": {
  "type": "string",
```

```

    "anyOf": [
      {
        "pattern": "1.0.0"
      }
    ]
  },
  "required": [
    "connectionConfiguration",
    "repositoryConfigurations",
    "additionalProperties",
    "type",
    "secretArn",
    "syncMode"
  ]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the data source.
repositoryEndpointMetadata	The endpoint information for the data source. This data source doesn't specify an endpoint in <code>repositoryEndpointMetadata</code> . Rather, the connection information is included in an AWS Secrets Manager secret that you provide the <code>secretArn</code> .
repositoryConfigurations	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings.
<ul style="list-style-type: none"> community user message attachment 	A list of objects that map attributes or field names of Microsoft Yammer objects to Amazon Q index field names.

Configuration	Description
secretARN	The Amazon Resource Name (ARN) of an AWS Secrets Manager secret that contains the key-value pairs required to connect to your Microsoft Yammer data source. This includes your client ID and client secret.
additionalProperties	Additional configuration options for your content in your data source
isCrawlAcl	Specify true to crawl access control information from documents.
fieldForUserId	Specify field to use for UserId for ACL crawling.
<ul style="list-style-type: none"> • isCrawlMessage • isCrawlAttachment • isCrawlPrivateMessage 	Input TRUE to index
<ul style="list-style-type: none"> • sinceDate 	Use to specify the time from when Amazon Q should crawl your Microsoft Yammer content
<ul style="list-style-type: none"> • communityNameFilter 	Use to specify community names to index.
inclusionPatterns	A list of regular expression patterns to <i>include</i> specific files in your Yammer data source. Files that match the patterns are included in the index. File that don't match the patterns are excluded from the index. Files that match the patterns are included in the index. Files that don't match the patterns are excluded from the index. If a file matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence and the file isn't included in the index.

Configuration	Description
<code>exclusionPatterns</code>	A list of regular expression patterns to <i>exclude</i> specific files in your Yammer data source. Files that match the patterns are excluded from the files in your data source. Files that match the patterns are excluded from the index. Files that don't match the patterns are included in the index. If a file matches both an exclusion and inclusion pattern, the exclusion pattern takes precedence and the file isn't included in the index.
<code>type</code>	Specify YAMMER as your data source type
<code>useChangeLog</code>	<code>true</code> to use the Yammer change log to determine which documents require adding, updating, or deleting in the index.
<code>syncMode</code>	Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose between the following options: <ul style="list-style-type: none">• Use <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index• Use <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index• Use <code>CHANGE_LOG</code> to incrementally crawl only new and modified content each time your data source syncs with your index

Configuration	Description
<code>enableIdentityCrawler</code>	true to activate identity crawler. Identity crawler is activated by default. Crawling identity information on users and groups with access to certain documents is useful for user context filtering. Search results are filtered based on the user or their group access to documents. See Identity crawler for more information.

How Amazon Q Business connector crawls Microsoft Yammer ACLs

When you connect an Microsoft Yammer data source to Amazon Q Business, Amazon Q Business crawls ACL information attached to a document (user and group information) from your Microsoft Yammer instance. If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

The group and user IDs are mapped as follows:

- `_email_id` – Your Microsoft email ID is an identifier that's necessary to configure each connector instance. Your email ID can be found in the properties section of your Microsoft account dashboard.
- `_group_id` – Group IDs exist in Microsoft Yammer Instances where there are set access permissions. They're mapped from the names of the groups in Microsoft Yammer.
- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business Microsoft Yammer data source connector field mappings

To help you structure data for retrieval and chat filtering, Amazon Q Business crawls data source document attributes or metadata and maps them to fields in your Amazon Q index.

Amazon Q has reserved fields that it uses when querying your application. When possible, Amazon Q automatically maps these built-in fields to attributes in your data source. If a built-in field doesn't have a default mapping, or if you want to map additional index fields, use the custom field mappings to specify how a data source attribute maps to your Amazon Q application. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q Yammer connector supports the following entities and the associated reserved and custom attributes.

Note

You can map any Yammer field to the document title or document body Amazon Q reserved/default index fields.

Supported entities and field mappings

- [Message](#)
- [Attachment](#)
- [User](#)
- [Community](#)

Message

Amazon Q supports crawling [Microsoft Yammer Messages](#) and offers the following message field mappings.

Microsoft Yammer field name	Index field name	Description	Data type
id	ymr_id	Custom	String
message_type	ymr_message_type	Custom	String
api_url	ymr_api_url	Custom	String
group_id	ymr_group_id	Custom	String
group_name	ymr_group_name	Custom	String
in_private_conversation	ymr_in_private_conversation	Custom	String
in_private_group	ymr_in_private_group	Custom	String
sender_email	ymr_sender_email	Custom	String
sender_id	ymr_sender_id	Custom	String
sender_name	ymr_sender_name	Custom	String
created_at	_created_at	Default	Date
web_url	_source_uri	Default	String

Attachment

Microsoft Yammer field name	Index field name	Description	Data type
id	ymr_attachment_id	Custom	String
name	ymr_attachment_name	Custom	String

Microsoft Yammer field name	Index field name	Description	Data type
size	ymr_attachment_size	Custom	String
url	ymr_attachment_url	Custom	String
file_type	ymr_attachment_type	Custom	String
created_at	_created_at	Default	Date
privacy	ymr_attachment_privacy	Custom	String
group_name	ymr_attachment_group_name	Custom	String
sender_email	ymr_attachment_sender_email	Custom	String
web_url	_source_uri	Default	String

User

Microsoft Yammer field name	Index field name	Description	Data type
id	ymr_user_id	Custom	String
user_type	ymr_user_type	Custom	String
state	ymr_user_state	Custom	String
full_name	ymr_user_full_name	Custom	String
activated_at	_created_at	Default	Date
first_name	ymr_user_first_name	Custom	String

Microsoft Yammer field name	Index field name	Description	Data type
last_name	ymr_user_last_name	Custom	String
network_name	ymr_user_network_name	Custom	String
network_domains	ymr_user_network_domains	Custom	String
url	ymr_user_url	Custom	String
name	ymr_user_name	Custom	String
birth_date	ymr_user_birth_date	Custom	Date
admin	ymr_user_admin	Custom	String
verified_admin	ymr_user_verified_admin	Custom	String
contact	ymr_user_contact	Custom	String
email	ymr_user_email	Custom	String
web_url	_source_uri	Default	String

Community

Microsoft Yammer field name	Index field name	Description	Data type
id	ymr_community_id	Custom	String
name	ymr_community_name	Custom	String

Microsoft Yammer field name	Index field name	Description	Data type
email	ymr_community_email	Custom	String
full_name	ymr_community_full_name	Custom	String
description	ymr_community_description	Custom	String
privacy	ymr_community_privacy	Custom	String
url	ymr_community_url	Custom	String
created_at	_created_at	Default	Date
state	ymr_community_state	Custom	String
web_url	_source_uri	Default	String

IAM role for Amazon Q Business Microsoft Yammer connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the BatchPutDocument and BatchDeleteDocument operations to ingest documents.

- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToGetSecret",
      "Effect": "Allow",
      "Action": [
        "secretsmanager:GetSecretValue"
      ],
      "Resource": [
        "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
      ]
    },
    {
      "Sid": "AllowsAmazonQToDecryptSecret",
      "Effect": "Allow",
      "Action": [
        "kms:Decrypt"
      ],
      "Resource": [
        "arn:aws:kms:{{region}}:{{account_id}}:key/[key_id]"
      ],
      "Condition": {
        "StringLike": {
          "kms:ViaService": [
            "secretsmanager.*.amazonaws.com"
          ]
        }
      }
    },
    {
      "Sid": "AllowsAmazonQToIngestDocuments",
      "Effect": "Allow",
      "Action": [
        "qbusiness:BatchPutDocument",
```

```

    "qbusiness:BatchDeleteDocument"
  ],
  "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
},
{
  "Sid": "AllowsAmazonQToIngestPrincipalMapping",
  "Effect": "Allow",
  "Action": [
    "qbusiness:PutGroup",
    "qbusiness:CreateUser",
    "qbusiness>DeleteGroup",
    "qbusiness:UpdateUser",
    "qbusiness:ListGroups"
  ],
  "Resource": [
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}",
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}",
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}/data-source/*"
  ]
},
{
  "Sid": "AllowsAmazonQToCreateAndDeleteNI",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateNetworkInterface",
    "ec2>DeleteNetworkInterface"
  ],
  "Resource": [
    "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
    "arn:aws:ec2:{{region}}:{{account_id}}:security-group/[{{security_group}}]"
  ]
},
{
  "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateNetworkInterface",
    "ec2>DeleteNetworkInterface"
  ],
  "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
  "Condition": {

```

```

    "StringLike": {
      "aws:RequestTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
    },
    "ForAllValues:StringEquals": {
      "aws:TagKeys": [
        "AMAZON_Q"
      ]
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateTags",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateTags"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringEquals": {
        "ec2:CreateAction": "CreateNetworkInterface"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterfacePermission"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:ResourceTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
    "Effect": "Allow",
    "Action": [
      "ec2:DescribeNetworkInterfaces",
      "ec2:DescribeAvailabilityZones",
      "ec2:DescribeNetworkInterfaceAttribute",
      "ec2:DescribeVpcs",

```

```

    "ec2:DescribeRegions",
    "ec2:DescribeNetworkInterfacePermissions",
    "ec2:DescribeSubnets"
  ],
  "Resource": "*"
}
]
}
```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "{{source_account}}"
        },
        "ArnEquals": {
          "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
        }
      }
    }
  ]
}
```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Known limitations for the Amazon Q Business Microsoft Yammer connector

The Amazon Q Business Microsoft Yammer connector has the following known limitations:

- Due to API limitations, an incremental sync will not update deleted **Messages**, **Attachments**, **Communities** and **Users**. To update deleted entities, you must run a full sync.

Troubleshooting your Amazon Q Business Microsoft Yammer connector

The following table provides information about error codes you may see for the Microsoft Yammer connector and suggested troubleshooting actions.

Error code	Error message	Suggested resolution
YMR-5001	Authentication error.	Provide a valid client id, client secret, username, password.
YMR-5002	Error validating credentials due to invalid username or password.	Provide a valid username and password.
YMR-5003	Error validating credentials due to invalid client id or client secret.	Provide valid client id and client secret.
YMR-5004	Access token is empty or null .	Provide non-empty or non-null access token .
YMR-5100	Null/empty client id.	Provide client id.
YMR-5101	Null/empty client secret.	Provide client secret.
YMR-5102	Null/empty username.	Provide username.
YMR-5103	Null/empty password.	Provide password .
YMR-5104	Null/empty since date.	Provide sinceDate .
YMR-5105	invalid sinceDate format.	since date format should be like YYYY-MM-DDTHH:mm:ss+00:00.
YMR-5106	Empty/null Repository configurations.	Provide Repository configurations.
YMR-5107	Empty/null message entity in repository configuration.	Provide message entity in repository configuration.

Error code	Error message	Suggested resolution
YMR-5108	Empty/null Attachment entity in repository configuration.	Provide attachment entity in repository configuration.
YMR-5109	Empty/null message entity field mapping.	Provide message entity field mapping.
YMR-5110	Empty/null attachment entity field mapping.	Provide attachment entity field mapping.
YMR-5111	Empty/null indexFieldName, indexFieldType or dataSourceFieldName in message entity.	Provide value for indexFieldName, indexFieldType and dataSourceFieldName in message entity.
YMR-5112	Empty/null indexFieldName, indexFieldType or dataSourceFieldName in attachment entity.	Provide value for indexFieldName, indexFieldType and dataSourceFieldName in message entity.
YMR-5113	Invalid patterns in the regex.	Provide valid regex patterns.
YMR-5114	Since date should be less than current date.	Provide since date less than current date.
YMR-5115	Only String, Date and Long formats are supported for field mappings.	Provide String, Date and Long formats for field mappings.
YMR-5116	Null/empty Network Domain.	Provide Network Domain.
YMR-5117	Got error while building groups.	Refer to logs for more information.

Error code	Error message	Suggested resolution
YMR-5118	Configuration found null during change access token.	Please provide valid configurations.
YMR-5119	Unable to connect to Yammer account.	Refer to logs for more details.
YMR-5120	An error occurred during the test connection.	Refer to logs for more details.
YMR-5500	Bad zip entry.	Provide a valid zip file.
YMR-5501	Invalid URI.	Provide valid URI.
YMR-5502	ContinuableInternalServerError.	Try again later.

Connecting MySQL to Amazon Q Business

MySQL is an open source relational database management system. You can connect your MySQL instance to Amazon Q Business—using either the AWS Management Console, CLI, or the [CreateDataSource](#) API—and create an Amazon Q web experience.

The Amazon Q MySQL data source connector supports MySQL 8.0. 21.

Important

As a best practice, provide Amazon Q with read-only database credentials. Also, avoid adding tables with sensitive data or personal identifiable information (PII).

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).

- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [MySQL connector overview](#)
- [Prerequisites for connecting Amazon Q Business to MySQL](#)
- [Connecting Amazon Q Business to MySQL using the console](#)
- [Connecting Amazon Q Business to MySQL using APIs](#)
- [How Amazon Q Business connector crawls MySQL ACLs](#)
- [Amazon Q Business MySQL data source connector field mappings](#)
- [IAM role for Amazon Q Business MySQL connector](#)
- [Known limitations for the Amazon Q Business MySQL connector](#)

MySQL connector overview

The following table gives an overview of the Amazon Q Business MySQL connector and its supported features.

Category	Feature	Support
Security	Authentication type	Basic
	Authentication credentials	<ul style="list-style-type: none"> • Username of database user • Password of database user
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Driver version	MySQL – 8.0.27
	Data source version	MySQL 8.0.21
	Identity crawling	No
	VPC	Yes

Category	Feature	Support
Crawl features	Custom metadata	Yes
	Entities	Yes. The following entities are supported: <ul style="list-style-type: none"> Document <div data-bbox="862 464 1508 730" style="border: 1px solid #add8e6; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p>Note</p> <p>Each database row is considered an individual searchable Amazon Q document.</p> </div>
	Field mappings	Yes. Supports both default and custom field mappings. For more information, see Field mappings .
	Sync mode	Supports full and incremental sync.
	File types	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to MySQL

Before you begin, make sure that you have completed the following prerequisites.

In MySQL, make sure you have:

- Noted your database user name and password.

Important

As a best practice, provide Amazon Q with read-only database credentials.

- Copied your database host URL, port, and instance.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your MySQL authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Connecting Amazon Q Business to MySQL using the console

The following procedure outlines how to connect Amazon Q Business to MySQL using the AWS Management Console.

Connecting Amazon Q to MySQL

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **MySQL** page, enter the following information:
6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. In **Source**, enter the following information:
 - a. **Host** – Enter the database host name.
 - b. **Port** – Enter the database port.

- c. **Instance** – Enter the database instance.
 - d. **Enable SSL certificate location** – Choose to enter the Amazon S3 path to your SSL certificate file.
8. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

 **Note**

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

9. In **Authentication** – Enter the following information for your **AWS Secrets Manager secret**.
- a. **Secret name** – A name for your secret.
 - b. For **Database user name**, and **Password** – Enter the authentication credential values you copied from your database.
 - c. Choose **Save**.
10. **Configure VPC and security group** – *optional* – Choose whether you want to use a VPC. If you do, enter the following information:
- a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
 - b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

11. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

12. In **Sync scope**, enter the following information:

- **SQL query** – Enter SQL query statements like SELECT and JOIN operations. SQL queries must be less than 1000 characters and not contain any semi-colons (;). Amazon Q will crawl all database content that matches your query.
- **Primary key column** – Provide the primary key for the database table. This identifies a table within your database.
- **Title column** – Provide the name of the document title column within your database table.
- **Body column** – Provide the name of the document body column within your database table.

13. In **Additional configuration – optional** – Configure the following settings:

- **Change-detecting columns** – Enter the names of the columns that Amazon Q will use to detect content changes. Amazon Q will re-index content when there is a change in any of these columns.
- **Users' IDs column** – Enter the name of the column which contains User IDs to be allowed access to content.
- **Groups column** – Enter the name of the column that contains groups to be allowed access to content.
- **Source URLs column** – Enter the name of the column which contains Source URLs to be indexed.
- **Time stamps column** – Enter the name of the column which contains time stamps. Amazon Q uses time stamp information to detect changes in your content and sync only changed content.
- **Time zones column** – Enter the name of the column which contains time zones for the content to be crawled.
- **Time stamps format** – Enter the name of the column which contains time stamp formats to use to detect content changes and re-sync your content.

14. In **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.

- **Full sync** – Sync all content regardless of the previous sync status.
- **New or modified content sync** – Sync only new and modified documents.
- **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.

For more details, see [Sync mode](#).

15. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
16. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
17. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:
 - a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
 - b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

 **Note**

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

18. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

19. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

 **Note**

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to MySQL using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the configuration parameter to provide a JSON schema with all other configuration information specific to your data source connector.

MySQL JSON schema

The following is the MySQL JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
      "type": "object",
      "properties": {
        "repositoryEndpointMetadata": {
          "type": "object",
          "properties": {
            "dbType": {
              "type": "string",
              "enum": [
                "mysql",
                "db2",
                "postgresql",
                "oracle",
                "sqlserver"
              ]
            },
            "dbHost": {
              "type": "string"
            },
            "dbPort": {
              "type": "string"
            },
            "dbInstance": {
              "type": "string"
            }
          }
        },
        "required": [
          "dbType",
          "dbHost",
```

```

        "dbPort",
        "dbInstance"
    ]
}
},
"required": [
    "repositoryEndpointMetadata"
]
},
"repositoryConfigurations": {
    "type": "object",
    "properties": {
        "document": {
            "type": "object",
            "properties": {
                "fieldMappings": {
                    "type": "array",
                    "items": [
                        {
                            "type": "object",
                            "properties": {
                                "indexFieldName": {
                                    "type": "string"
                                },
                                "indexFieldType": {
                                    "type": "string"
                                },
                                "dataSourceFieldName": {
                                    "type": "string"
                                }
                            }
                        },
                        {
                            "type": "object",
                            "properties": {
                                "indexFieldName": {
                                    "type": "string"
                                },
                                "indexFieldType": {
                                    "type": "string"
                                },
                                "dataSourceFieldName": {
                                    "type": "string"
                                }
                            }
                        }
                    ]
                }
            }
        }
    },
    "required": [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
    ]
}
}
},
"required": [
    "fieldMappings"
]
}
}

```

```
  },
  "required": [
  ]
},
"additionalProperties": {
  "type": "object",
  "properties": {
    "primaryKey": {
      "type": "string"
    },
    "titleColumn": {
      "type": "string"
    },
    "bodyColumn": {
      "type": "string"
    },
    "sqlQuery": {
      "type": "string",
      "not": {
        "pattern": ";+"
      }
    },
    "timestampColumn": {
      "type": "string"
    },
    "timestampFormat": {
      "type": "string"
    },
    "timezone": {
      "type": "string"
    },
    "changeDetectingColumns": {
      "type": "array",
      "items": {
        "type": "string"
      }
    },
    "allowedUsersColumn": {
      "type": "string"
    },
    "allowedGroupsColumn": {
      "type": "string"
    },
    "sourceURIColumn": {
```

```

        "type": "string"
    },
    "serverlessAurora": {
        "type": "string",
        "enum": ["true", "false"]
    }
},
"required": ["primaryKey", "titleColumn", "bodyColumn", "sqlQuery"]
},
"type" : {
    "type" : "string",
    "pattern": "JDBC"
},
"syncMode": {
    "type": "string",
    "enum": [
        "FORCED_FULL_CRAWL",
        "FULL_CRAWL"
    ]
},
"secretArn": {
    "type": "string"
}
},
"version": {
    "type": "string",
    "anyOf": [
        {
            "pattern": "1.0.0"
        }
    ]
}
},
"required": [
    "connectionConfiguration",
    "repositoryConfigurations",
    "syncMode",
    "additionalProperties",
    "secretArn",
    "type"
]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.
repositoryEndpointMetadata	<p>Required configuration information for connecting your data source.</p> <ul style="list-style-type: none"> • dbType—The type of Java database you are using, whether <code>mysql</code>, <code>db2</code>, <code>postgresql</code>, <code>oracle</code>, or <code>sqlserver</code>. • dbHost—The database host name. • dbPort—The database port. • dbInstance—The database instance.
repositoryConfigurations	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings. Specify the type of data source and the secret ARN.
document	A list of objects that map the attributes or field names of your database content to Amazon Q index field names. For more information, see Mapping data source fields .
additionalProperties	Additional configuration options for your content in your data source. Use to include or exclude specific content in your database data source.
primaryKey	Provide the primary key for the database table. This identifies a table within your database.
titleColumn	Provide the name of the document title column within your database table.

Configuration	Description
bodyColumn	Provide the name of the document title column within your database table.
sqlQuery	Enter SQL query statements like SELECT and JOIN operations. SQL queries must be less than 1000 characters and not contain any semi-colons (;). Amazon Q will crawl all database content that matches your query.
timestampColumn	Enter the name of the column which contains time stamps. Amazon Q uses time stamp information to detect changes in your content and sync only changed content.
timestampFormat	Enter the name of the column which contains time stamp formats to use to detect content changes and re-sync your content.
timezone	Enter the name of the column which contains time zones for the content to be crawled.
changeDetectingColumns	Enter the names of the columns that Amazon Q will use to detect content changes. Amazon Q will re-index content when there is a change in any of these columns
allowedUsersColumns	Enter the name of the column which contains User IDs to be allowed access to content.
allowedGroupsColumn	Enter the name of the column which contains User IDs to be allowed access to content.
sourceURIColumn	Enter the name of the column which contains Source URLs to be indexed.
isSslEnabled	true to add a path to an SSL certificate file stored in an Amazon S3 bucket.

Configuration	Description
type	The type of data source. Specify JDBC as your data source type.
syncMode	<p>Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose</p> <ul style="list-style-type: none">• <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index• <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index• <code>CHANGE_LOG</code> to incrementally crawl only new and modified content each time your data source syncs with your index.
secretArn	<p>The Amazon Resource Name (ARN) of a Secrets Manager secret that contains user name and password required to connect to your database. The secret must contain a JSON structure with the following keys:</p> <pre data-bbox="829 1373 1507 1570">{ "user name": "<i>database user name</i>", "password": "<i>password</i>" }</pre>
version	The version of the template that is currently supported.

How Amazon Q Business connector crawls MySQL ACLs

When you connect a database data source to Amazon Q, Amazon Q crawls user and group information from a column in the source table. You specify this column in the console or using the configuration parameter as part of the `CreateDataSource` operation.

If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

A database data source has the following limitations:

- You can only specify an allow list for a database data source. You can't specify a deny list.
- You can only specify groups. You can't specify individual users for the allow list.
- The database column should be a string containing a semicolon delimited list of groups.

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business MySQL data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a attribute mapping already available, or if you want to map additional

document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q PostgreSQL connector supports the following field mappings:

Supported field mappings

- [Document](#)

Document

JDBC field name	Index field name	Description	Data type
jd_document_id	jd_document_id	Custom	String
jd_document_title	jd_document_title	Custom	String
jd_source_uri	_source_uri	Default	String

IAM role for Amazon Q Business MySQL connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the BatchPutDocument and BatchDeleteDocument operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- Permission to access the SSL certificate stored in your Amazon S3 bucket.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AllowsAmazonQToGetS3Objects",
    "Action": [
      "s3:GetObject"
    ],
    "Resource": [
      "arn:aws:s3:::{{input_bucket_name}}/*"
    ],
    "Effect": "Allow",
    "Condition": {
      "StringEquals": {
        "aws:ResourceAccount": "{{account_id}}"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToGetSecret",
    "Effect": "Allow",
    "Action": [
      "secretsmanager:GetSecretValue"
    ],
    "Resource": [
      "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToDecryptSecret",
    "Effect": "Allow",
    "Action": [
```

```

        "kms:Decrypt"
    ],
    "Resource": [
        "arn:aws:kms:{{region}}:{{account_id}}:key/[key_id]"
    ],
    "Condition": {
        "StringLike": {
            "kms:ViaService": [
                "secretsmanager.*.amazonaws.com"
            ]
        }
    }
},
{
    "Sid": "AllowsAmazonQToIngestDocuments",
    "Effect": "Allow",
    "Action": [
        "qbusiness:BatchPutDocument",
        "qbusiness:BatchDeleteDocument"
    ],
    "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
},
{
    "Sid": "AllowsAmazonQToIngestPrincipalMapping",
    "Effect": "Allow",
    "Action": [
        "qbusiness:PutGroup",
        "qbusiness:CreateUser",
        "qbusiness>DeleteGroup",
        "qbusiness:UpdateUser",
        "qbusiness:ListGroups"
    ],
    "Resource": [
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}",
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}",
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}/data-source/*"
    ]
},
{
    "Sid": "AllowsAmazonQToCreateAndDeleteNI",

```

```

    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2>DeleteNetworkInterface"
    ],
    "Resource": [
      "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
      "arn:aws:ec2:{{region}}:{{account_id}}:security-group/
[[security_group]]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2>DeleteNetworkInterface"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:RequestTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
      },
      "ForAllValues:StringEquals": {
        "aws:TagKeys": [
          "AMAZON_Q"
        ]
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateTags",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateTags"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringEquals": {
        "ec2:CreateAction": "CreateNetworkInterface"
      }
    }
  }
},

```

```

    {
      "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
      "Effect": "Allow",
      "Action": [
        "ec2:CreateNetworkInterfacePermission"
      ],
      "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
      "Condition": {
        "StringLike": {
          "aws:ResourceTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
        }
      }
    },
    {
      "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
      "Effect": "Allow",
      "Action": [
        "ec2:DescribeNetworkInterfaces",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeNetworkInterfaceAttribute",
        "ec2:DescribeVpcs",
        "ec2:DescribeRegions",
        "ec2:DescribeNetworkInterfacePermissions",
        "ec2:DescribeSubnets"
      ],
      "Resource": "*"
    }
  ]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToAssumeRoleForServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
    }
  ]
}

```

```
    "Condition": {
      "StringEquals": {
        "aws:SourceAccount": "{{source_account}}"
      },
      "ArnLike": {
        "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
      }
    }
  }
]
```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Known limitations for the Amazon Q Business MySQL connector

- Deleted database rows will not be tracked in when Amazon Q checks for updated content.
- The size of field names and values in a row of your database can't exceed 400KB.
- If you have a large amount of data in your database data source, and do not want Amazon Q to index all your database content after the first sync, you can choose to sync only new, modified, or deleted documents.

Connecting Oracle Database to Amazon Q Business

Oracle Database is a database management system. You can connect your Oracle Database instance to Amazon Q Business—using either the AWS Management Console, CLI, or the [CreateDataSource](#) API—and create an Amazon Q web experience.

The Amazon Q Oracle Database data source connector supports Oracle Database 18c, 19c, and 21c.

Important

As a best practice, provide Amazon Q with read-only database credentials. Also, avoid adding tables with sensitive data or personal identifiable information (PII).

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).
- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [Oracle Database connector overview](#)
- [Prerequisites for connecting Amazon Q Business to Oracle Database](#)
- [Connecting Amazon Q Business to Oracle Database using the console](#)
- [Connecting Amazon Q Business to Oracle Database using APIs](#)
- [How Amazon Q Business connector crawls Oracle Database ACLs](#)
- [Amazon Q Business Oracle Database data source connector field mappings](#)
- [IAM role for Amazon Q Business Oracle Database connector](#)
- [Known limitations for the Amazon Q Business Oracle Database connector](#)

Oracle Database connector overview

The following table gives an overview of the Amazon Q Oracle Database connector and its supported features.

Category	Feature	Support
Security	Authentication type	Basic
	Authentication credentials	<ul style="list-style-type: none"> • Username of database user • Password of database user
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Driver version	Oracle – 21.1.0.0
	Data source version	Oracle Database 18c, 19c, 21c

Category	Feature	Support
	Identity crawling	No
	VPC	Yes
Crawl features	Custom metadata	Yes
	Entities	Yes. The following entities are supported: <ul style="list-style-type: none"> Document <div data-bbox="862 625 1511 894" style="border: 1px solid #add8e6; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p> Note Each database row is considered an individual searchable Amazon Q document.</p> </div>
	Field mappings	Yes. Supports both default and custom field mappings. For more information, see Field mappings .
	Sync mode	Supports full and incremental sync.
	File types	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to Oracle Database

Before you begin, make sure that you have completed the following prerequisites.

In Oracle Database, make sure you have:

- Noted your database user name and password.

Important

As a best practice, provide Amazon Q with read-only database credentials.

- Copied your database host URL, port, and instance.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your Oracle Database authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Connecting Amazon Q Business to Oracle Database using the console

The following procedure outlines how to connect Amazon Q to Oracle Database using the AWS Management Console.

Connecting Amazon Q to Oracle Database

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **Oracle Database** page, enter the following information:
6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. In **Source**, enter the following information:
 - a. **Host** – Enter the database host name.

- b. **Port** – Enter the database port.
 - c. **Instance** – Enter the database instance.
 - d. **Enable SSL certificate location** – Choose to enter the Amazon S3 path to your SSL certificate file.
8. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

 **Note**

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

9. In **Authentication** – Enter the following information for your **AWS Secrets Manager secret**.
- a. **Secret name** – A name for your secret.
 - b. For **Database user name**, and **Password** – Enter the authentication credential values you copied from your database.
 - c. Choose **Save**.
10. **Configure VPC and security group** – *optional* – Choose whether you want to use a VPC. If you do, enter the following information:
- a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
 - b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

11. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

12. In **Sync scope**, enter the following information:

- **SQL query** – Enter SQL query statements like SELECT and JOIN operations. SQL queries must be less than 1000 characters and not contain any semi-colons (;). Amazon Q will crawl all database content that matches your query.
- **Primary key column** – Provide the primary key for the database table. This identifies a table within your database.
- **Title column** – Provide the name of the document title column within your database table.
- **Body column** – Provide the name of the document body column within your database table.

13. In **Additional configuration – optional** – Configure the following settings:

- **Change-detecting columns** – Enter the names of the columns that Amazon Q will use to detect content changes. Amazon Q will re-index content when there is a change in any of these columns.
- **Users' IDs column** – Enter the name of the column which contains User IDs to be allowed access to content.
- **Groups column** – Enter the name of the column that contains groups to be allowed access to content.
- **Source URLs column** – Enter the name of the column which contains Source URLs to be indexed.
- **Time stamps column** – Enter the name of the column which contains time stamps. Amazon Q uses time stamp information to detect changes in your content and sync only changed content.
- **Time zones column** – Enter the name of the column which contains time zones for the content to be crawled.
- **Time stamps format** – Enter the name of the column which contains time stamp formats to use to detect content changes and re-sync your content.

14. In **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.

- **Full sync** – Sync all content regardless of the previous sync status.
- **New or modified content sync** – Sync only new and modified documents.
- **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.

For more details, see [Sync mode](#).

15. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
16. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
17. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:
 - a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
 - b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

 **Note**

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

18. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

19. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

 **Note**

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to Oracle Database using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the configuration parameter to provide a JSON schema with all other configuration information specific to your data source connector.

Oracle Database JSON schema

The following is the Oracle Database JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
      "type": "object",
      "properties": {
        "repositoryEndpointMetadata": {
          "type": "object",
          "properties": {
            "dbType": {
              "type": "string",
              "enum": [
                "mysql",
                "db2",
                "postgresql",
                "oracle",
                "sqlserver"
              ]
            },
            "dbHost": {
              "type": "string"
            },
            "dbPort": {
              "type": "string"
            },
            "dbInstance": {
              "type": "string"
            }
          }
        },
        "required": [
          "dbType",
          "dbHost",
```

```

        "dbPort",
        "dbInstance"
    ]
}
},
"required": [
    "repositoryEndpointMetadata"
]
},
"repositoryConfigurations": {
    "type": "object",
    "properties": {
        "document": {
            "type": "object",
            "properties": {
                "fieldMappings": {
                    "type": "array",
                    "items": [
                        {
                            "type": "object",
                            "properties": {
                                "indexFieldName": {
                                    "type": "string"
                                },
                                "indexFieldType": {
                                    "type": "string"
                                },
                                "dataSourceFieldName": {
                                    "type": "string"
                                }
                            }
                        },
                        {
                            "type": "object",
                            "properties": {
                                "indexFieldName": {
                                    "type": "string"
                                },
                                "indexFieldType": {
                                    "type": "string"
                                },
                                "dataSourceFieldName": {
                                    "type": "string"
                                }
                            }
                        }
                    ]
                }
            }
        }
    },
    "required": [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
    ]
}
}
},
"required": [
    "fieldMappings"
]
}
}

```

```
    },
    "required": [
    ]
  },
  "additionalProperties": {
    "type": "object",
    "properties": {
      "primaryKey": {
        "type": "string"
      },
      "titleColumn": {
        "type": "string"
      },
      "bodyColumn": {
        "type": "string"
      },
      "sqlQuery": {
        "type": "string",
        "not": {
          "pattern": ";+"
        }
      },
      "timestampColumn": {
        "type": "string"
      },
      "timestampFormat": {
        "type": "string"
      },
      "timezone": {
        "type": "string"
      },
      "changeDetectingColumns": {
        "type": "array",
        "items": {
          "type": "string"
        }
      },
      "allowedUsersColumn": {
        "type": "string"
      },
      "allowedGroupsColumn": {
        "type": "string"
      },
      "sourceURIColumn": {
```

```

        "type": "string"
    },
    "serverlessAurora": {
        "type": "string",
        "enum": ["true", "false"]
    }
},
"required": ["primaryKey", "titleColumn", "bodyColumn", "sqlQuery"]
},
"type" : {
    "type" : "string",
    "pattern": "JDBC"
},
"syncMode": {
    "type": "string",
    "enum": [
        "FORCED_FULL_CRAWL",
        "FULL_CRAWL"
    ]
},
"secretArn": {
    "type": "string"
}
},
"version": {
    "type": "string",
    "anyOf": [
        {
            "pattern": "1.0.0"
        }
    ]
}
},
"required": [
    "connectionConfiguration",
    "repositoryConfigurations",
    "syncMode",
    "additionalProperties",
    "secretArn",
    "type"
]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.
repositoryEndpointMetadata	<p>Required configuration information for connecting your data source.</p> <ul style="list-style-type: none"> • dbType—The type of Java database you are using, whether <code>mysql</code>, <code>db2</code>, <code>postgresql</code>, <code>oracle</code>, or <code>sqlserver</code>. • dbHost—The database host name. • dbPort—The database port. • dbInstance—The database instance.
repositoryConfigurations	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings. Specify the type of data source and the secret ARN.
document	A list of objects that map the attributes or field names of your database content to Amazon Q index field names. For more information, see Mapping data source fields .
additionalProperties	Additional configuration options for your content in your data source. Use to include or exclude specific content in your database data source.
primaryKey	Provide the primary key for the database table. This identifies a table within your database.
titleColumn	Provide the name of the document title column within your database table.

Configuration	Description
bodyColumn	Provide the name of the document title column within your database table.
sqlQuery	Enter SQL query statements like SELECT and JOIN operations. SQL queries must be less than 1000 characters and not contain any semi-colons (;). Amazon Q will crawl all database content that matches your query.
timestampColumn	Enter the name of the column which contains time stamps. Amazon Q uses time stamp information to detect changes in your content and sync only changed content.
timestampFormat	Enter the name of the column which contains time stamp formats to use to detect content changes and re-sync your content.
timezone	Enter the name of the column which contains time zones for the content to be crawled.
changeDetectingColumns	Enter the names of the columns that Amazon Q will use to detect content changes. Amazon Q will re-index content when there is a change in any of these columns
allowedUsersColumns	Enter the name of the column which contains User IDs to be allowed access to content.
allowedGroupsColumn	Enter the name of the column which contains User IDs to be allowed access to content.
sourceURIColumn	Enter the name of the column which contains Source URLs to be indexed.
isSslEnabled	true to add a path to an SSL certificate file stored in an Amazon S3 bucket.

Configuration	Description
type	The type of data source. Specify JDBC as your data source type.
syncMode	<p>Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose</p> <ul style="list-style-type: none">• <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index• <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index• <code>CHANGE_LOG</code> to incrementally crawl only new and modified content each time your data source syncs with your index.
secretArn	<p>The Amazon Resource Name (ARN) of a Secrets Manager secret that contains user name and password required to connect to your database. The secret must contain a JSON structure with the following keys:</p> <pre data-bbox="829 1373 1507 1570">{ "user name": "<i>database user name</i>", "password": "<i>password</i>" }</pre>
version	The version of the template that is currently supported.

How Amazon Q Business connector crawls Oracle Database ACLs

When you connect a database data source to Amazon Q, Amazon Q crawls user and group information from a column in the source table. You specify this column in the console or using the configuration parameter as part of the `CreateDataSource` operation.

If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

A database data source has the following limitations:

- You can only specify an allow list for a database data source. You can't specify a deny list.
- You can only specify groups. You can't specify individual users for the allow list.
- The database column should be a string containing a semicolon delimited list of groups.

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business Oracle Database data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a attribute mapping already available, or if you want to map additional

document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q PostgreSQL connector supports the following field mappings:

Supported field mappings

- [Document](#)

Document

JDBC field name	Index field name	Description	Data type
jd_document_id	jd_document_id	Custom	String
jd_document_title	jd_document_title	Custom	String
jd_source_uri	_source_uri	Default	String

IAM role for Amazon Q Business Oracle Database connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the BatchPutDocument and BatchDeleteDocument operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- Permission to access the SSL certificate stored in your Amazon S3 bucket.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AllowsAmazonQToGetS3Objects",
    "Action": [
      "s3:GetObject"
    ],
    "Resource": [
      "arn:aws:s3:::{{input_bucket_name}}/*"
    ],
    "Effect": "Allow",
    "Condition": {
      "StringEquals": {
        "aws:ResourceAccount": "{{account_id}}"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToGetSecret",
    "Effect": "Allow",
    "Action": [
      "secretsmanager:GetSecretValue"
    ],
    "Resource": [
      "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToDecryptSecret",
    "Effect": "Allow",
    "Action": [
```

```

        "kms:Decrypt"
    ],
    "Resource": [
        "arn:aws:kms:{{region}}:{{account_id}}:key/[key_id]"
    ],
    "Condition": {
        "StringLike": {
            "kms:ViaService": [
                "secretsmanager.*.amazonaws.com"
            ]
        }
    }
},
{
    "Sid": "AllowsAmazonQToIngestDocuments",
    "Effect": "Allow",
    "Action": [
        "qbusiness:BatchPutDocument",
        "qbusiness:BatchDeleteDocument"
    ],
    "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
},
{
    "Sid": "AllowsAmazonQToIngestPrincipalMapping",
    "Effect": "Allow",
    "Action": [
        "qbusiness:PutGroup",
        "qbusiness:CreateUser",
        "qbusiness>DeleteGroup",
        "qbusiness:UpdateUser",
        "qbusiness:ListGroup"
    ],
    "Resource": [
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}",
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}",
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}/data-source/*"
    ]
},
{
    "Sid": "AllowsAmazonQToCreateAndDeleteNI",

```

```

    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2>DeleteNetworkInterface"
    ],
    "Resource": [
      "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
      "arn:aws:ec2:{{region}}:{{account_id}}:security-group/
[[security_group]]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2>DeleteNetworkInterface"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:RequestTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
      },
      "ForAllValues:StringEquals": {
        "aws:TagKeys": [
          "AMAZON_Q"
        ]
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateTags",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateTags"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringEquals": {
        "ec2:CreateAction": "CreateNetworkInterface"
      }
    }
  }
},

```



```

    {
      "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
      "Effect": "Allow",
      "Action": [
        "ec2:CreateNetworkInterfacePermission"
      ],
      "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
      "Condition": {
        "StringLike": {
          "aws:ResourceTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
        }
      }
    },
    {
      "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
      "Effect": "Allow",
      "Action": [
        "ec2:DescribeNetworkInterfaces",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeNetworkInterfaceAttribute",
        "ec2:DescribeVpcs",
        "ec2:DescribeRegions",
        "ec2:DescribeNetworkInterfacePermissions",
        "ec2:DescribeSubnets"
      ],
      "Resource": "*"
    }
  ]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToAssumeRoleForServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
    }
  ]
}

```

```
    "Condition": {
      "StringEquals": {
        "aws:SourceAccount": "{{source_account}}"
      },
      "ArnLike": {
        "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
      }
    }
  }
]
```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Known limitations for the Amazon Q Business Oracle Database connector

- Deleted database rows will not be tracked in when Amazon Q checks for updated content.
- The size of field names and values in a row of your database can't exceed 400KB.
- If you have a large amount of data in your database data source, and do not want Amazon Q to index all your database content after the first sync, you can choose to sync only new, modified, or deleted documents.

Connecting PostgreSQL to Amazon Q Business

PostgreSQL is an open source database management system. You can connect your PostgreSQL instance to Amazon Q Business—using either the AWS Management Console, CLI, or the [CreateDataSource](#) API—and create an Amazon Q web experience.

The Amazon Q PostgreSQL data source connector supports PostgreSQL 9.6.

Important

As a best practice, provide Amazon Q with read-only database credentials. Also, avoid adding tables with sensitive data or personal identifiable information (PII).

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).
- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [PostgreSQL connector overview](#)
- [Prerequisites for connecting Amazon Q Business to PostgreSQL](#)
- [Connecting Amazon Q Business to PostgreSQL using the console](#)
- [Connecting Amazon Q Business to PostgreSQL using APIs](#)
- [How Amazon Q Business connector crawls PostgreSQL ACLs](#)
- [Amazon Q Business PostgreSQL data source connector field mappings](#)
- [IAM role for Amazon Q Business PostgreSQL connector](#)
- [Known limitations for the Amazon Q Business PostgreSQL connector](#)

PostgreSQL connector overview

The following table gives an overview of the Amazon Q Business PostgreSQL connector and its supported features.

Category	Feature	Support
Security	Authentication type	Basic
	Authentication credentials	<ul style="list-style-type: none"> • Username of database user • Password of database user
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Driver version	PostgreSQL – 42.3.2
	Data source version	PostgreSQL 9.6

Category	Feature	Support
	Identity crawling	No
	VPC	Yes
Crawl features	Custom metadata	Yes
	Entities	Yes. The following entities are supported: <ul style="list-style-type: none"> Document <div data-bbox="862 625 1507 892" style="border: 1px solid #add8e6; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p>Note</p> <p>Each database row is considered an individual searchable Amazon Q document.</p> </div>
	Field mappings	Yes. Supports both default and custom field mappings. For more information, see Field mappings .
	Sync mode	Supports full and incremental sync.
	File types	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to PostgreSQL

Before you begin, make sure that you have completed the following prerequisites.

In PostgreSQL, make sure you have:

- Noted your database user name and password.

Important

As a best practice, provide Amazon Q with read-only database credentials.

- Copied your database host URL, port, and instance.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your PostgreSQL authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Connecting Amazon Q Business to PostgreSQL using the console

The following procedure outlines how to connect Amazon Q Business to PostgreSQL using the AWS Management Console.

Connecting Amazon Q to PostgreSQL

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **PostgreSQL** page, enter the following information:
6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. In **Source**, enter the following information:
 - a. **Host** – Enter the database host URL.

- b. **Port** – Enter the database port, for example, 5432.
 - c. **Instance** – Enter the database instance, for example postgres.
 - d. **Enable SSL certificate location** – Choose to enter the Amazon S3 path to your SSL certificate file.
8. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

 **Note**

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

9. In **Authentication** – Enter the following information for your **AWS Secrets Manager secret**.
 - a. **Secret name** – A name for your secret.
 - b. For **Database user name**, and **Password** – Enter the authentication credential values you copied from your database.
 - c. Choose **Save**.
10. **Configure VPC and security group** – *optional* – Choose whether you want to use a VPC. If you do, enter the following information:
 - a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
 - b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

11. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

12. In **Sync scope**, enter the following information:

- **SQL query** – Enter SQL query statements like SELECT and JOIN operations. SQL queries must be less than 1000 characters and not contain any semi-colons (;). Amazon Q will crawl all database content that matches your query.
- **Primary key column** – Provide the primary key for the database table. This identifies a table within your database.
- **Title column** – Provide the name of the document title column within your database table.
- **Body column** – Provide the name of the document body column within your database table.

13. In **Additional configuration – optional** – Configure the following settings:

- **Change-detecting columns** – Enter the names of the columns that Amazon Q will use to detect content changes. Amazon Q will re-index content when there is a change in any of these columns.
- **Users' IDs column** – Enter the name of the column which contains User IDs to be allowed access to content.
- **Groups column** – Enter the name of the column that contains groups to be allowed access to content.
- **Source URLs column** – Enter the name of the column which contains Source URLs to be indexed.
- **Time stamps column** – Enter the name of the column which contains time stamps. Amazon Q uses time stamp information to detect changes in your content and sync only changed content.
- **Time zones column** – Enter the name of the column which contains time zones for the content to be crawled.
- **Time stamps format** – Enter the name of the column which contains time stamp formats to use to detect content changes and re-sync your content.

14. In **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.

- **Full sync** – Sync all content regardless of the previous sync status.
- **New or modified content sync** – Sync only new and modified documents.
- **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.

For more details, see [Sync mode](#).

15. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
16. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
17. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:
 - a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
 - b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

 **Note**

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

18. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

19. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

 **Note**

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to PostgreSQL using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the configuration parameter to provide a JSON schema with all other configuration information specific to your data source connector.

PostgreSQL JSON schema

The following is the PostgreSQL JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
      "type": "object",
      "properties": {
        "repositoryEndpointMetadata": {
          "type": "object",
          "properties": {
            "dbType": {
              "type": "string",
              "enum": [
                "mysql",
                "db2",
                "postgresql",
                "oracle",
                "sqlserver"
              ]
            },
            "dbHost": {
              "type": "string"
            },
            "dbPort": {
              "type": "string"
            },
            "dbInstance": {
              "type": "string"
            }
          }
        },
        "required": [
          "dbType",
          "dbHost",
```

```
        "dbPort",
        "dbInstance"
    ]
}
},
"required": [
    "repositoryEndpointMetadata"
]
},
"repositoryConfigurations": {
    "type": "object",
    "properties": {
        "document": {
            "type": "object",
            "properties": {
                "fieldMappings": {
                    "type": "array",
                    "items": [
                        {
                            "type": "object",
                            "properties": {
                                "indexFieldName": {
                                    "type": "string"
                                },
                                "indexFieldType": {
                                    "type": "string"
                                },
                                "dataSourceFieldName": {
                                    "type": "string"
                                }
                            }
                        },
                        {
                            "type": "object",
                            "properties": {
                                "indexFieldName": {
                                    "type": "string"
                                },
                                "indexFieldType": {
                                    "type": "string"
                                },
                                "dataSourceFieldName": {
                                    "type": "string"
                                }
                            }
                        }
                    ]
                }
            }
        }
    },
    "required": [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
    ]
}
}
},
"required": [
    "fieldMappings"
]
}
```

```
    },
    "required": [
    ]
  },
  "additionalProperties": {
    "type": "object",
    "properties": {
      "primaryKey": {
        "type": "string"
      },
      "titleColumn": {
        "type": "string"
      },
      "bodyColumn": {
        "type": "string"
      },
      "sqlQuery": {
        "type": "string",
        "not": {
          "pattern": ";+"
        }
      },
      "timestampColumn": {
        "type": "string"
      },
      "timestampFormat": {
        "type": "string"
      },
      "timezone": {
        "type": "string"
      },
      "changeDetectingColumns": {
        "type": "array",
        "items": {
          "type": "string"
        }
      },
      "allowedUsersColumn": {
        "type": "string"
      },
      "allowedGroupsColumn": {
        "type": "string"
      },
      "sourceURIColumn": {
```

```

        "type": "string"
    },
    "serverlessAurora": {
        "type": "string",
        "enum": ["true", "false"]
    }
},
"required": ["primaryKey", "titleColumn", "bodyColumn", "sqlQuery"]
},
"type" : {
    "type" : "string",
    "pattern": "JDBC"
},
"syncMode": {
    "type": "string",
    "enum": [
        "FORCED_FULL_CRAWL",
        "FULL_CRAWL"
    ]
},
"secretArn": {
    "type": "string"
}
},
"version": {
    "type": "string",
    "anyOf": [
        {
            "pattern": "1.0.0"
        }
    ]
}
},
"required": [
    "connectionConfiguration",
    "repositoryConfigurations",
    "syncMode",
    "additionalProperties",
    "secretArn",
    "type"
]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.
repositoryEndpointMetadata	<p>Required configuration information for connecting your data source.</p> <ul style="list-style-type: none"> • dbType—The type of Java database you are using, whether <code>mysql</code>, <code>db2</code>, <code>postgresql</code>, <code>oracle</code>, or <code>sqlserver</code>. • dbHost—The database host name. • dbPort—The database port. • dbInstance—The database instance.
repositoryConfigurations	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings. Specify the type of data source and the secret ARN.
document	A list of objects that map the attributes or field names of your database content to Amazon Q index field names. For more information, see Mapping data source fields .
additionalProperties	Additional configuration options for your content in your data source. Use to include or exclude specific content in your database data source.
primaryKey	Provide the primary key for the database table. This identifies a table within your database.
titleColumn	Provide the name of the document title column within your database table.

Configuration	Description
bodyColumn	Provide the name of the document title column within your database table.
sqlQuery	Enter SQL query statements like SELECT and JOIN operations. SQL queries must be less than 1000 characters and not contain any semi-colons (;). Amazon Q will crawl all database content that matches your query.
timestampColumn	Enter the name of the column which contains time stamps. Amazon Q uses time stamp information to detect changes in your content and sync only changed content.
timestampFormat	Enter the name of the column which contains time stamp formats to use to detect content changes and re-sync your content.
timezone	Enter the name of the column which contains time zones for the content to be crawled.
changeDetectingColumns	Enter the names of the columns that Amazon Q will use to detect content changes. Amazon Q will re-index content when there is a change in any of these columns
allowedUsersColumns	Enter the name of the column which contains User IDs to be allowed access to content.
allowedGroupsColumn	Enter the name of the column which contains User IDs to be allowed access to content.
sourceURIColumn	Enter the name of the column which contains Source URLs to be indexed.
isSslEnabled	true to add a path to an SSL certificate file stored in an Amazon S3 bucket.

Configuration	Description
type	The type of data source. Specify JDBC as your data source type.
syncMode	<p>Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose</p> <ul style="list-style-type: none">• <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index• <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index• <code>CHANGE_LOG</code> to incrementally crawl only new and modified content each time your data source syncs with your index.
secretArn	<p>The Amazon Resource Name (ARN) of a Secrets Manager secret that contains user name and password required to connect to your database. The secret must contain a JSON structure with the following keys:</p> <pre data-bbox="829 1373 1507 1570">{ "user name": "<i>database user name</i>", "password": "<i>password</i>" }</pre>
version	The version of the template that is currently supported.

How Amazon Q Business connector crawls PostgreSQL ACLs

When you connect a database data source to Amazon Q, Amazon Q crawls user and group information from a column in the source table. You specify this column in the console or using the configuration parameter as part of the `CreateDataSource` operation.

If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

A database data source has the following limitations:

- You can only specify an allow list for a database data source. You can't specify a deny list.
- You can only specify groups. You can't specify individual users for the allow list.
- The database column should be a string containing a semicolon delimited list of groups.

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business PostgreSQL data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a attribute mapping already available, or if you want to map additional

document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q PostgreSQL connector supports the following field mappings:

Supported field mappings

- [Document](#)

Document

JDBC field name	Index field name	Description	Data type
jd_document_id	jd_document_id	Custom	String
jd_document_title	jd_document_title	Custom	String
jd_source_uri	_source_uri	Default	String

IAM role for Amazon Q Business PostgreSQL connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the BatchPutDocument and BatchDeleteDocument operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- Permission to access the SSL certificate stored in your Amazon S3 bucket.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AllowsAmazonQToGetS3Objects",
    "Action": [
      "s3:GetObject"
    ],
    "Resource": [
      "arn:aws:s3:::{{input_bucket_name}}/*"
    ],
    "Effect": "Allow",
    "Condition": {
      "StringEquals": {
        "aws:ResourceAccount": "{{account_id}}"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToGetSecret",
    "Effect": "Allow",
    "Action": [
      "secretsmanager:GetSecretValue"
    ],
    "Resource": [
      "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToDecryptSecret",
    "Effect": "Allow",
    "Action": [
```

```

        "kms:Decrypt"
    ],
    "Resource": [
        "arn:aws:kms:{{region}}:{{account_id}}:key/[key_id]"
    ],
    "Condition": {
        "StringLike": {
            "kms:ViaService": [
                "secretsmanager.*.amazonaws.com"
            ]
        }
    }
},
{
    "Sid": "AllowsAmazonQToIngestDocuments",
    "Effect": "Allow",
    "Action": [
        "qbusiness:BatchPutDocument",
        "qbusiness:BatchDeleteDocument"
    ],
    "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
},
{
    "Sid": "AllowsAmazonQToIngestPrincipalMapping",
    "Effect": "Allow",
    "Action": [
        "qbusiness:PutGroup",
        "qbusiness:CreateUser",
        "qbusiness>DeleteGroup",
        "qbusiness:UpdateUser",
        "qbusiness:ListGroup"
    ],
    "Resource": [
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}",
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}",
        "arn:aws:qbusiness:{{region}}:{{account_id}}:application/
{{application_id}}/index/{{index_id}}/data-source/*"
    ]
},
{
    "Sid": "AllowsAmazonQToCreateAndDeleteNI",

```

```

    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2>DeleteNetworkInterface"
    ],
    "Resource": [
      "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
      "arn:aws:ec2:{{region}}:{{account_id}}:security-group/
[[security_group]]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2>DeleteNetworkInterface"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:RequestTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
      },
      "ForAllValues:StringEquals": {
        "aws:TagKeys": [
          "AMAZON_Q"
        ]
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateTags",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateTags"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringEquals": {
        "ec2:CreateAction": "CreateNetworkInterface"
      }
    }
  },

```

```

    {
      "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
      "Effect": "Allow",
      "Action": [
        "ec2:CreateNetworkInterfacePermission"
      ],
      "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
      "Condition": {
        "StringLike": {
          "aws:ResourceTag/AMAZON_Q":
"qbusiness_{{account_id}}_{{application_id}}_*"
        }
      }
    },
    {
      "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
      "Effect": "Allow",
      "Action": [
        "ec2:DescribeNetworkInterfaces",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeNetworkInterfaceAttribute",
        "ec2:DescribeVpcs",
        "ec2:DescribeRegions",
        "ec2:DescribeNetworkInterfacePermissions",
        "ec2:DescribeSubnets"
      ],
      "Resource": "*"
    }
  ]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToAssumeRoleForServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
    }
  ]
}

```

```
    "Condition": {
      "StringEquals": {
        "aws:SourceAccount": "{{source_account}}"
      },
      "ArnLike": {
        "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
      }
    }
  }
]
```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Known limitations for the Amazon Q Business PostgreSQL connector

- Deleted database rows will not be tracked in when Amazon Q checks for updated content.
- The size of field names and values in a row of your database can't exceed 400KB.
- If you have a large amount of data in your database data source, and do not want Amazon Q to index all your database content after the first sync, you can choose to sync only new, modified, or deleted documents.

Connecting Quip to Amazon Q Business

Quip is a collaborative productivity software that offers real time document-authoring capabilities. You can connect your Quip instance to Amazon Q Business—using either the AWS Management Console, CLI, or the [CreateDataSource](#) API—and create an Amazon Q web experience.

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).
- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [Quip connector overview](#)
- [Prerequisites for connecting Amazon Q Business to Quip](#)
- [Retrieving Quip credentials](#)
- [Connecting Amazon Q Business to Quip using the console](#)
- [Connecting Amazon Q Business to Quip using APIs](#)
- [How Amazon Q Business connector crawls Quip ACLs](#)
- [Amazon Q Business Quip data source connector field mappings](#)
- [IAM role for Amazon Q Business Quip connector](#)
- [Known limitations for the Amazon Q Business Quip connector](#)

Quip connector overview

The following table gives an overview of the Amazon Q Business Quip connector and its supported features.

Category	Feature	Support
Security	Authentication type	Personal Access Token
	Authentication credentials	Quip token
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Identity crawling	No. Quip doesn't have the concept of user groups.
	VPC	Yes
Crawl features	Custom metadata	No
	Entities	Yes. The following entities are supported: <ul style="list-style-type: none"> • Thread • Message • Attachment

Category	Feature	Support
	Field mappings	Yes. Supports default field mappings. For more information, see Field mappings .
	Filters	Yes. The following filters are supported: <ul style="list-style-type: none"> • Crawl file comments • Crawl chat rooms • Crawl attachments
	Sync mode	Supports full and incremental (new, modified, and deleted) sync.
	File types	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to Quip

Before you begin, make sure that you have completed the following prerequisites.

In Quip, make sure you have:

- A Quip account with administrative permissions.
- Created Quip authentication credentials that include a personal access token. See [Quip documentation on authentication](#) for more information.
- Copied your Quip site domain. For example, *https://quip-company.quipdomain.com/browse* where *quipdomain* is the domain.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your Quip authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Retrieving Quip credentials

Before you connect Quip to Amazon Q, you need to create and retrieve the Quip credentials you will use to connect Quip to Amazon Q.

The following procedure gives you an overview of how to configure Quip for connecting with Amazon Q by creating a API access token.

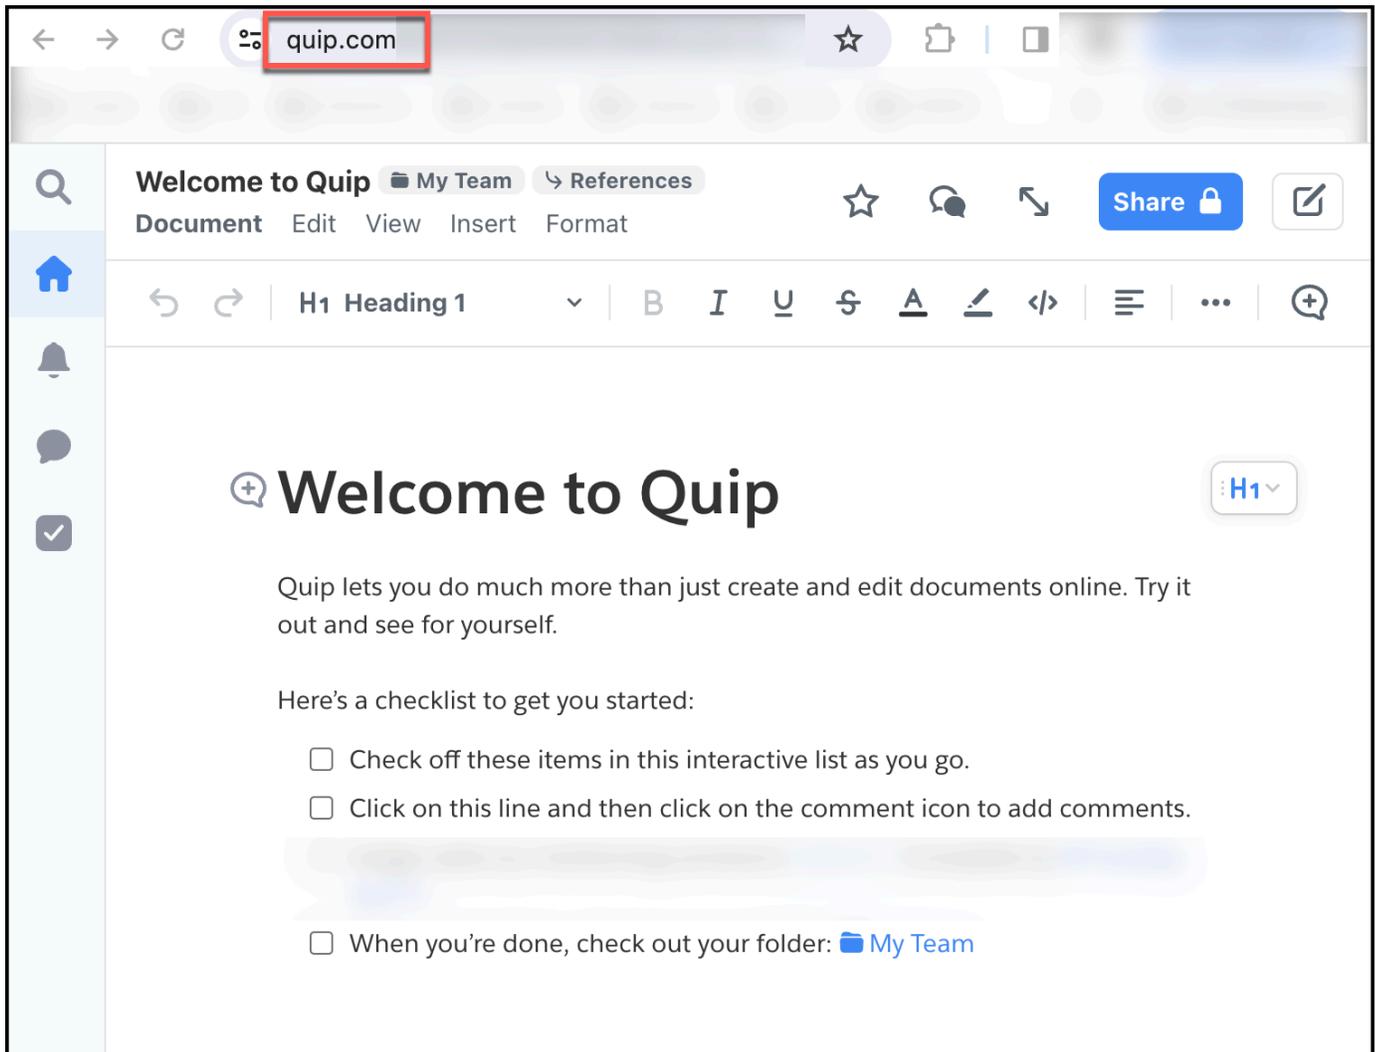
Configuring Quip authentication for Amazon Q

1. Log in to your Quip account using a web browser of your choice and sign into your Quip workspace.

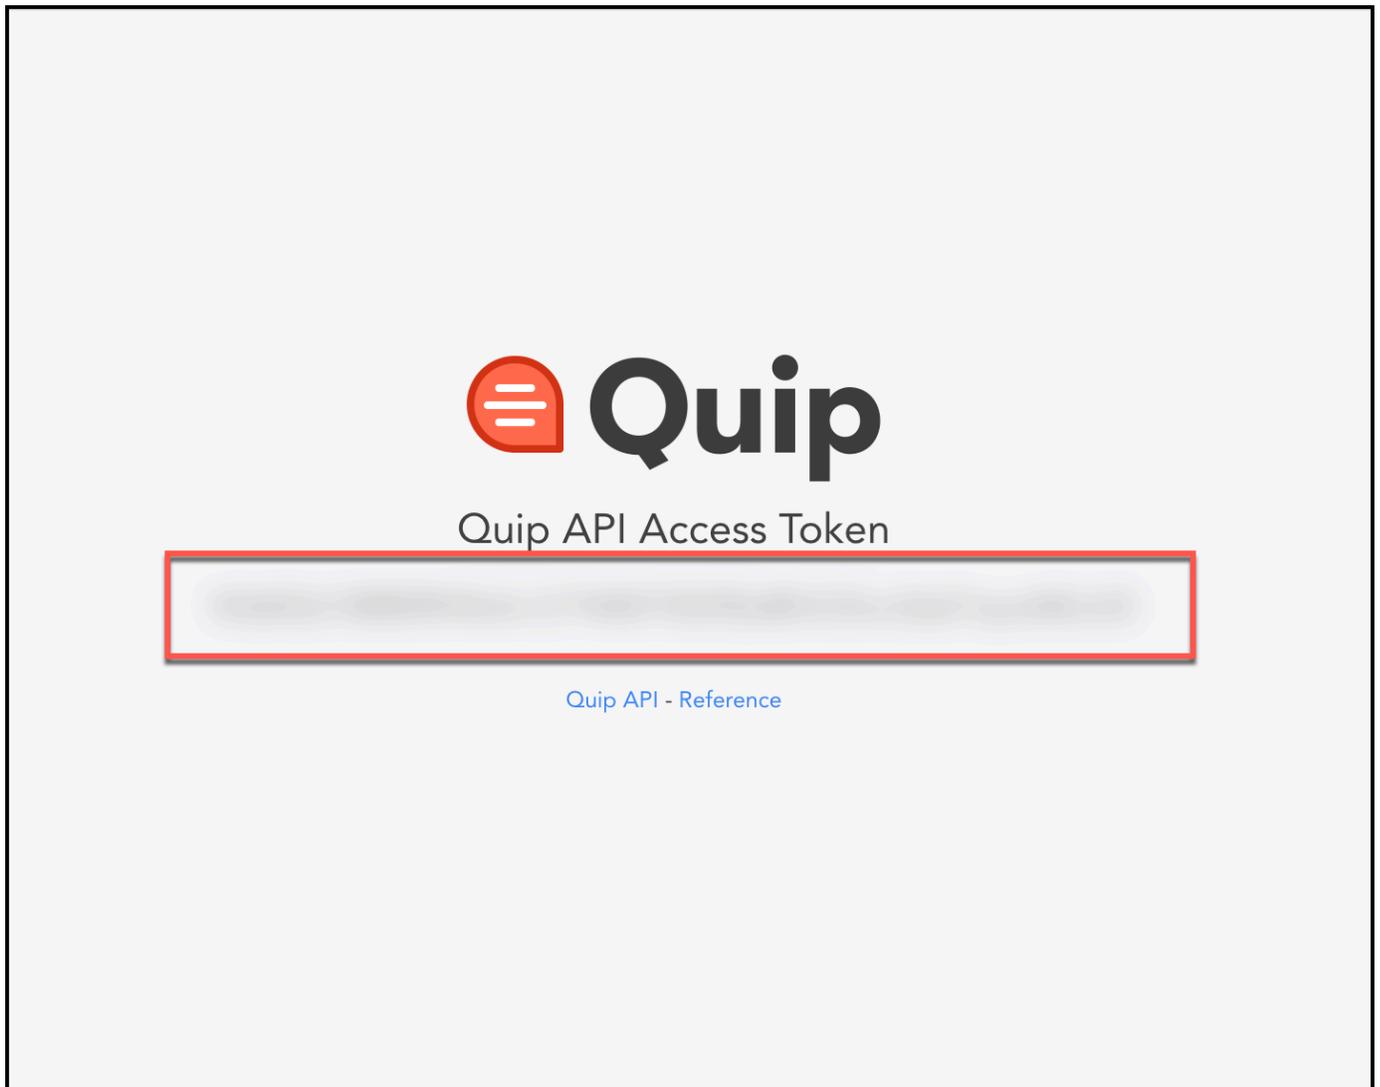
Note

To configure Quip for Amazon Q, you must be an admin user in the Quip account.

2. From the browser URL, note your Quip domain name. You will need this both to connect to Amazon Q and also to generate an API access token.



3. In a text editor of your choice, copy and paste the following: `https://domain/dev/token`. Then, replace *domain* with the Quip domain you copied in the last step. Copy the URL.
4. Open a new browser window and paste the formatted URL you created in the last step. Quip will return an API access token in your browser window.



You now have the Quip domain name and Quip API access token you need to connect to Amazon Q.

Connecting Amazon Q Business to Quip using the console

The following procedure outlines how to connect Amazon Q Business to Quip using the AWS Management Console.

Connecting Amazon Q to Quip

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).

3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **Quip** page, enter the following information:
6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. **Source** – Enter your **Quip domain name**. You can find your domain name in the browser URL of your Quip. For example, *https://quip-company.quipdomain.com/browse*, the domain is "quipdomain".
8. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

 **Note**

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

9. **Authentication** – Enter the following information for your **AWS Secrets Manager secret**.
 - a. **Secret name** – A name for your secret.
 - b. **Quip token** – Enter the Quip personal access token you created in your Quip account.
10. **Configure VPC and security group** – *optional* – Choose whether you want to use a VPC. If you do, enter the following information:
 - a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
 - b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

11. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

12. In **Sync scope**, enter the following information:

- a. **Add Quip folder IDs to crawl** – Enter the Quip folder IDs you want to crawl. You can find your folder ID in the browser URL when you access your folder in Quip. For example, <https://quip-company.quipdomain.com/zlLu0VNSarTL/folder-name>, the folder ID is "zlLu0VNSarTL"..

 **Note**

To crawl a root folder, including all sub-folders and documents inside it, input the root folder ID. To crawl specific sub-folders, add the specific sub-folder IDs.

- b. (Optional) **Additional configuration – optional** – Configure the following settings:
 - **Content types** – Choose between crawling **All content**, **File comments**, **Chat rooms** and **Attachments**.
 - **Regex patterns** – Add regex patterns to include or exclude file names, file types, or file paths. You can have a total of 100 patterns.
13. For **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.
 - **Full sync** – Sync all content regardless of the previous sync status.
 - **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.
 14. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
 15. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
 16. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:

- a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
- b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

 **Note**

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

17. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

18. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

 **Note**

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to Quip using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the `configuration` parameter to provide a JSON schema with all other configuration information specific to your data source connector.

For an example of the API request, see [CreateDataSource](#) in the Amazon Q API Reference.

Quip JSON schema

The following is the Quip JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
      "type": "object",
      "properties": {
        "repositoryEndpointMetadata": {
          "type": "object",
          "properties": {
            "domain": {
              "type": "string"
            }
          },
          "required": [
            "domain"
          ]
        }
      },
      "required": [
        "repositoryEndpointMetadata"
      ]
    },
    "repositoryConfigurations": {
      "type": "object",
      "properties": {
        "thread": {
          "type": "object",
          "properties": {
            "fieldMappings": {
              "type": "array",
              "items": [
                {
                  "type": "object",
                  "properties": {
                    "indexFieldName": {
                      "type": "string"
                    },
                    "indexFieldType": {
                      "type": "string",

```

```

        "enum": [
            "STRING",
            "STRING_LIST",
            "DATE"
        ]
    },
    "dataSourceFieldName": {
        "type": "string"
    },
    "dateFieldFormat": {
        "type": "string",
        "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
    }
},
"required": [
    "indexFieldName",
    "indexFieldType",
    "dataSourceFieldName"
]
}
]
}
},
"required": [
    "fieldMappings"
]
},
"message": {
    "type": "object",
    "properties": {
        "fieldMappings": {
            "type": "array",
            "items": [
                {
                    "type": "object",
                    "properties": {
                        "indexFieldName": {
                            "type": "string"
                        },
                    },
                    "indexFieldType": {
                        "type": "string",
                        "enum": [
                            "STRING",
                            "STRING_LIST",

```



```

        "DATE"
      ]
    },
    "dataSourceFieldName": {
      "type": "string"
    },
    "dateFieldFormat": {
      "type": "string",
      "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
    }
  },
  "required": [
    "indexFieldName",
    "indexFieldType",
    "dataSourceFieldName"
  ]
}
]
}
},
"required": [
  "fieldMappings"
]
},
"attachment": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": [
        {
          "type": "object",
          "properties": {
            "indexFieldName": {
              "type": "string"
            },
            "indexFieldType": {
              "type": "string",
              "enum": [
                "STRING",
                "STRING_LIST",
                "DATE"
              ]
            }
          }
        }
      ]
    }
  }
},

```

```

        "dataSourceFieldName": {
            "type": "string"
        },
        "dateFieldFormat": {
            "type": "string",
            "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
        }
    },
    "required": [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
    ]
}
]
}
},
"required": [
    "fieldMappings"
]
}
},
"additionalProperties": {
    "type": "object",
    "properties": {
        "isCrawlAcl": {
            "type": "boolean"
        },
        "fieldForUserId": {
            "type": "string"
        },
        "folderIds": {
            "type": "array",
            "items": {
                "type": "string"
            }
        },
        "crawlFileComments": {
            "type": "boolean"
        },
        "crawlChatRooms": {
            "type": "boolean"
        }
    }
},

```

```
        "crawlAttachments": {
            "type": "boolean"
        },
        "inclusionPatterns": {
            "type": "array",
            "items": {
                "type": "string"
            }
        },
        "exclusionPatterns": {
            "type": "array",
            "items": {
                "type": "string"
            }
        }
    },
    "required": []
},
"type": {
    "type": "string",
    "pattern": "QUIP"
},
"syncMode": {
    "type": "string",
    "enum": [
        "FULL_CRAWL",
        "FORCED_FULL_CRAWL"
    ]
},
"secretArn": {
    "type": "string",
    "minLength": 20,
    "maxLength": 2048
}
},
"version": {
    "type": "string",
    "anyOf": [
        {
            "pattern": "1.0.0"
        }
    ]
}
},
"required": [
```

```

    "connectionConfiguration",
    "repositoryConfigurations",
    "syncMode",
    "additionalProperties",
    "secretArn",
    "type"
  ]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.
repositoryEndpointMetadata	The endpoint information for the data source.
domain	Your Quip site domain. For example, <i>https://quip-company.quipdomain.com/browse</i> where <i>quipdomain</i> is the domain.
repositoryConfigurations	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings.
<ul style="list-style-type: none"> thread message attachment 	A list of objects that map the attributes or field names of your Quip pages and assets to Amazon Q index field names.
additionalProperties	Additional configuration options for your content in your data source.
isCrawlAcl	Specify true to crawl access control information from documents.
fieldForUserId	Specify field to use for UserId for ACL crawling.

Configuration	Description
folderIds	Specify folder IDs to crawl.
<ul style="list-style-type: none"> • crawlFileComments • crawlChatRooms • crawlAttachments 	true to index.
<ul style="list-style-type: none"> • inclusionPatterns 	A list of regular expression patterns to include specific content in your Quip data source. Content that matches the patterns are included in the index. Content that doesn't match the pattern are excluded from the index. If any content matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence, and the content isn't included in the index.
<ul style="list-style-type: none"> • exclusionPatterns 	A list of regular expression patterns to exclude specific content in your Quip data source. Content that matches the patterns are excluded from the index. Content that doesn't match the patterns are included in the index. If any content matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence, and the content isn't included in the index.
type	The type of data source. Specify QUIP as your data source type.
enableIdentityCrawler	Specify true to use the Amazon Q identity crawler to sync identity/principal information on users and groups with access to specific documents.

Configuration	Description
syncMode	<p>Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose between the following options:</p> <ul style="list-style-type: none"> • Use <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index. • Use <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index.
secretArn	<p>The Amazon Resource Name (ARN) of an AWS Secrets Manager secret that contains the key-value pairs required to connect to your Quip. The secret must contain a JSON structure with the following keys:</p> <pre data-bbox="829 1136 1507 1297"> { "accessToken": " <i>token</i>" } </pre>
version	<p>The version of this template that's currently supported.</p>

How Amazon Q Business connector crawls Quip ACLs

When you connect an Quip data source to Amazon Q Business, Amazon Q Business crawls ACL information attached to a document (user and group information) from your Quip instance. If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

The Quip user IDs are mapped as follows:

- `_user_id`—User IDs exist in Quip on files where there are set access permissions. They are mapped from the user emails as the IDs in Quip.

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business Quip data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q Business enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a attribute mapping already available, or if you want to map additional document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q Quip connector supports the following entities and the associated reserved and custom attributes.

Note

You can map any Quip field to the document title or document body Amazon Q reserved/default index fields.

Supported entities and field mappings

- [Thread](#)
- [Message](#)
- [Attachment](#)

Thread

Amazon Q supports crawling [Quip Threads](#) and offers the following thread field mappings.

Quip field name	Index field name	Description	Data type
qp_authors	_authors	Default	String list
qp_category	_category	Default	String
qp_file_type	qp_file_type	Custom	String
qp_document_title	qp_document_title	Custom	String
qp_source_uri	_source_uri	Default	String
qp_created_at	_created_at	Default	Date
qp_updated_at	_last_updated_at	Default	Date

Message

Amazon Q supports crawling [Quip Messages](#) and offers the following message field mappings.

Quip field name	Index field name	Description	Data type
qp_authors	_authors	Default	String list
qp_category	_category	Default	String
qp_source_uri	_source_uri	Default	String
qp_parent_file	qp_parent_file	Custom	String
qp_created_at	_created_at	Default	Date
qp_updated_at	_last_updated_at	Default	Date

Attachment

Amazon Q supports crawling Quip attachments and offers the following attachment field mappings.

Quip field name	Index field name	Description	Data type
qp_authors	_authors	Default	String list
qp_category	_category	Default	String
qp_source_uri	_source_uri	Default	String
qp_file_type	qp_file_type	Custom	String
qp_parent_file	qp_parent_file	Custom	String
qp_blob_id	qp_blob_id	Custom	String
qp_created_at	_created_at	Default	Date
qp_updated_at	_last_updated_at	Default	Date

IAM role for Amazon Q Business Quip connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the `BatchPutDocument` and `BatchDeleteDocument` operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToGetSecret",
      "Effect": "Allow",
      "Action": [
        "secretsmanager:GetSecretValue"
      ],
      "Resource": [
        "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
      ]
    },
    {
      "Sid": "AllowsAmazonQToDecryptSecret",
      "Effect": "Allow",
      "Action": [
        "kms:Decrypt"
      ],
      "Resource": [
```

```

    "arn:aws:kms:{{region}}:{{account_id}}:key/[{{key_id}}]"
  ],
  "Condition": {
    "StringLike": {
      "kms:ViaService": [
        "secretsmanager.*.amazonaws.com"
      ]
    }
  }
},
{
  "Sid": "AllowsAmazonQToIngestDocuments",
  "Effect": "Allow",
  "Action": [
    "qbusiness:BatchPutDocument",
    "qbusiness:BatchDeleteDocument"
  ],
  "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
},
{
  "Sid": "AllowsAmazonQToIngestPrincipalMapping",
  "Effect": "Allow",
  "Action": [
    "qbusiness:PutGroup",
    "qbusiness:CreateUser",
    "qbusiness>DeleteGroup",
    "qbusiness:UpdateUser",
    "qbusiness:ListGroups"
  ],
  "Resource": [
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}",
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}",
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}/data-source/*"
  ]
},
{
  "Sid": "AllowsAmazonQToCreateAndDeleteNI",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateNetworkInterface",
    "ec2>DeleteNetworkInterface"
  ]
}

```

```

    ],
    "Resource": [
      "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
      "arn:aws:ec2:{{region}}:{{account_id}}:security-group/[{{security_group}}]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2:DeleteNetworkInterface"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:RequestTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
      },
      "ForAllValues:StringEquals": {
        "aws:TagKeys": [
          "AMAZON_Q"
        ]
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateTags",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateTags"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringEquals": {
        "ec2:CreateAction": "CreateNetworkInterface"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterfacePermission"
    ],
  },

```

```

"Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
"Condition": {
  "StringLike": {
    "aws:ResourceTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
  }
},
{
  "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
  "Effect": "Allow",
  "Action": [
    "ec2:DescribeNetworkInterfaces",
    "ec2:DescribeAvailabilityZones",
    "ec2:DescribeNetworkInterfaceAttribute",
    "ec2:DescribeVpcs",
    "ec2:DescribeRegions",
    "ec2:DescribeNetworkInterfacePermissions",
    "ec2:DescribeSubnets"
  ],
  "Resource": "*"
}
]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "{{source_account}}"
        },
        "ArnEquals": {
          "aws:SourceArn": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/{{application_id}}"
        }
      }
    }
  ]
}

```

```
    }  
  }  
} ]  
}
```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Known limitations for the Amazon Q Business Quip connector

The Amazon Q Business Quip connector has the following known limitations:

- Only **Full sync** is supported by default. For **New, modified, or deleted content sync**, Admin API access is required and Admin API has to be enabled on the Quip website .
- Only data in shared folders will be crawled by the Amazon Q Quip connector. Private folders, other than the private folders belonging to the Private Access Token user, will not be crawled.
- Quip doesn't store file types and file paths. Amazon Q Quip connector can't support inclusion and exclusion filters on these.

Connecting Salesforce Online to Amazon Q Business

Salesforce is a customer relationship management (CRM) tool for managing support, sales, and marketing teams. You can connect Salesforce Online instance to Amazon Q Business—using either the AWS Management Console or the [CreateDataSource](#) API—and create an Amazon Q web experience.

The Amazon Q Salesforce Online connector supports the following Salesforce Online editions: Developer Edition and Enterprise Edition.

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).
- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [Salesforce Online connector overview](#)
- [Prerequisites for connecting Amazon Q Business to Salesforce Online](#)
- [Setting up Salesforce Online for connecting to Amazon Q Business](#)
- [Connecting Amazon Q Business to Salesforce Online using the console](#)
- [Connecting Amazon Q Business to Salesforce using APIs](#)
- [How Amazon Q Business connector crawls Salesforce ACLs](#)
- [Amazon Q Business Salesforce Online data source connector field mappings](#)
- [IAM role for Amazon Q Business Salesforce Online connector](#)
- [Known limitations for the Amazon Q Business Salesforce Online connector](#)
- [Troubleshooting your Amazon Q Business Salesforce Online connector](#)

Salesforce Online connector overview

The following table gives an overview of the Amazon Q Business Salesforce Online connector and its supported features.

Category	Feature	Support
Security	Authentication type	OAuth 2.0
	Authentication credentials	<ul style="list-style-type: none"> • Salesforce authentication URL • Username Client secret • Password username • Security token • Consumer key • Consumer secret
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Identity crawling	Yes
	VPC	Yes

Category	Feature	Support
	Supported versions	<ul style="list-style-type: none"> • API 30-56 • Lightning, Classic • Sandbox
Crawl features	Custom objects	Yes
	Custom metadata	Yes
	Entities	<p>Yes. The following entities are supported:</p> <ul style="list-style-type: none"> • Account • Campaign • Partner • Pricebook • Case • Contact • Contract • Document • Group • Idea • Lead • Opportunity • Product • Profile • Solution • Task • User • Chatter • Knowledge Articles

Category	Feature	Support
	Field mappings	Yes. Supports both default and custom field mappings. For more information, see Field mappings .
	Filters	Yes. The following filters are supported: <ul style="list-style-type: none"> • Attachment filter for supported entities • Regex filters for entities • Inclusion and exclusion filters on file type for Documents • Inclusion and exclusion filters on File Name and File Type for Attachments
	Sync mode	Supports full and incremental sync.
	File types	Supports all files supported by Amazon Q.

Category	Feature	Support
	Crawled as a document	<ul style="list-style-type: none"> • Each account • Each contact • Each campaign • Each contract • Each case • Each partner • Each opportunity • Each group • Each lead • Each user • Each task • Each idea • Each profile • Each solution • Each chatter • Each document • Each custom entity • Each knowledge article

Prerequisites for connecting Amazon Q Business to Salesforce Online

Before you begin, make sure that you have completed the following prerequisites.

In Salesforce, make sure you have:

- Copied the Salesforce security token associated with the account that's used to connect to Salesforce.
- Created a Salesforce Connected App account with OAuth activated and have copied the consumer key (client ID) and consumer secret (client secret) assigned to your Salesforce Connected App. For more information, see [Salesforce documentation on Connected Apps](#) on the Salesforce website.

- Copied the URL of the Salesforce instance that you want to index. Typically, this is <https://<company>.salesforce.com/>. The server must be running a Salesforce connected app.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your Salesforce Online authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Setting up Salesforce Online for connecting to Amazon Q Business

Before you connect Salesforce Online to Amazon Q Business, you need to create and retrieve the Salesforce Online credentials you will use to connect Salesforce Online to Amazon Q. You will also need to add any authorization permissions needed by Salesforce Online to connect to Amazon Q.

The following procedure gives you an overview of how to configure Salesforce Online for Amazon Q.

Configuring Salesforce Online for Amazon Q

1. Create a Salesforce Online instance at <https://developer.salesforce.com/signup>. Note the username and password you logged in with. Also note the Salesforce Online URL that's sent to your email on successful instance setup.

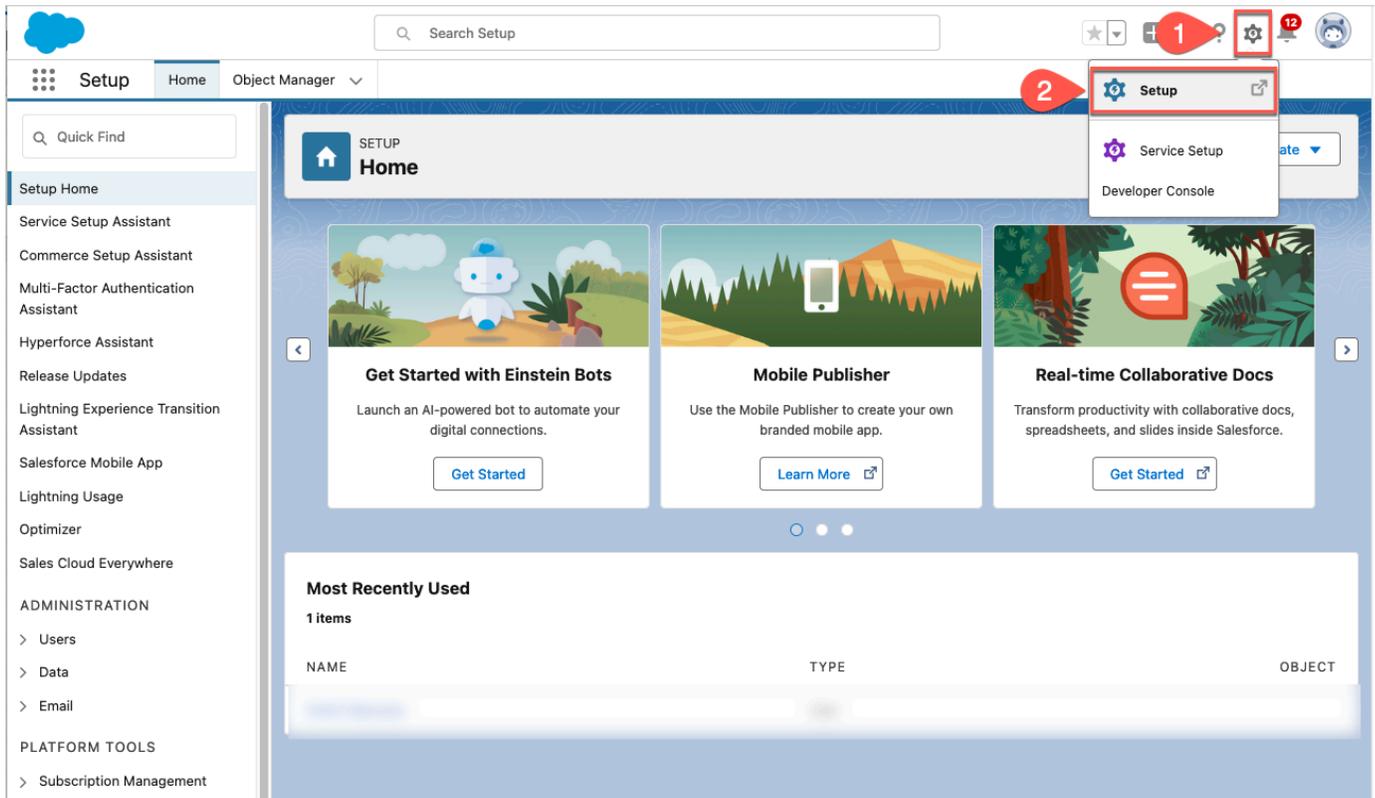
You will need these pieces of information later to connect to Amazon Q.

2. Log in to your Salesforce Online account at <https://login.salesforce.com>.
3. From the Salesforce Online profile menu, copy your Salesforce Online URL, if you haven't already. This will be the URL you will input as host URL in Amazon Q.

The screenshot shows the Salesforce Setup Home page. The top navigation bar includes a search bar labeled "Search Setup" and a user profile icon with a red notification badge containing the number "1". A red box highlights the profile icon, and a red circle with the number "2" points to the "Setup" option in the profile menu. The profile menu is open, showing options for "Settings" and "Log Out". Below the menu, the user's name and "DISPLAY DENSITY" (set to "Comfy") are visible. The main content area features three cards: "Get Started with Einstein Bots", "Mobile Publisher", and "Real-time Collaborative Docs". A "Most Recently Used" section is partially visible at the bottom, showing a table with columns for NAME, TYPE, and OBJECT.

NAME	TYPE	OBJECT

- Then, from the Salesforce Online profile menu, select the **Setup** icon and then select **Setup**.



5. From the left navigation menu, on the **Setup** home page, go to **Platform tools**, select **Apps**, and then, select **App manager**.

Then, from the **Lightning Experience App Manager** page, select **New Connected App**.

The screenshot shows the Salesforce Setup interface for the Lightning Experience App Manager. The top navigation bar includes the 'Setup' tab and a 'Home' tab (highlighted with a red box and callout 1). A search bar is located to the right of the navigation bar. The left sidebar contains a 'Quick Find' search box and a list of navigation items. The 'PLATFORM TOOLS' section is highlighted with a red box and callout 2. Under 'PLATFORM TOOLS', the 'Apps' item is highlighted with a red box and callout 3, and the 'App Manager' item is highlighted with a red box and callout 4. The main content area displays a table of 22 items, sorted by App Name. The table has columns for App Name, Developer Name, Description, Last Modified, and other details. A 'New Connected App' button is highlighted with a red box and callout 5.

6. On the **New Connected App** page, do the following:

- In **Basic information**, enter the following required information:
 - **Connect App Name** – A name for your connected app.
 - **API Name** – A name for your API.
 - **Contact Email** – Your contact email.

Basic Information |= Required Information

Connected App Name

API Name

Contact Email

Contact Phone

Logo Image URL
[Upload logo image](#) or [Choose one of our sample logos](#)

Icon URL
[Choose one of our sample logos](#)

Info URL

Description

Enter other values as per your use case.

- In **API (Enable OAuth Settings)**, select the checkbox to enable. Then, enter the following information:
 - **Callback URL** – Enter the following callback URL: `https://login.salesforce.com/services/oauth2/token`. Also, copy and save this URL in a text editor of your choice. You will enter this callback URL in Amazon Q later as **Authentication URL**.
 - **Select OAuth Scopes** – Select **Full access (full)** as your OAuth Scope.
 - **Introspect All Tokens** – Select this option to generate access tokens in a future step. You need this access token to connect to Amazon Q. You enter this as the **Security token** in the Amazon Q console.

API (Enable OAuth Settings)

Enable OAuth Settings 1

Enable for Device Flow

Callback URL 2

Use digital signatures

Selected OAuth Scopes

Available OAuth Scopes

- Access Chatbot services (chatbot_api)
- Access content resources (content)
- Access custom permissions (custom_permissions)
- Access the Salesforce API Platform (sfap_api)
- Access the identity URL service (id, profile, email, address, phone)
- Access unique user identifiers (openid)
- Manage Data Cloud Calculated Insight data (cdp_calculated_insight_api)
- Manage Data Cloud Identity Resolution (cdp_identityresolution_api)
- Manage Data Cloud Ingestion API data (cdp_ingest_api)
- Manage Data Cloud profile data (cdp_profile_api)
- Manage Pardot services (pardot_api)

Selected OAuth Scopes

Full access (full) 3

Require Proof Key for Code Exchange (PKCE) Extension for Supported Authorization Flows

Require Secret for Web Server Flow

Require Secret for Refresh Token Flow

Enable Client Credentials Flow

Enable Authorization Code and Credentials Flow

Enable Token Exchange Flow

Enable Refresh Token Rotation

Issue JSON Web Token (JWT)-based access tokens for named users

Introspect All Tokens 4

Configure ID Token

Enable Asset Tokens

Enable Single Logout

Select other options as per your use case.

- Select **Save**.
7. From the **Manage Connected Apps** page that opens, choose **Manage Consumer Details**. You will be redirected to a **Connected App Name** summary page.

SETUP
Manage Connected Apps

Connected App Name Help for this Page ?

• Back to List: Custom Apps

[Edit](#) [Delete](#) [Manage](#)

Changes can take up to 10 minutes to take effect. Deleting a parent org also deletes all connected apps with OAuth settings enabled.

Version
API Name
Created Date
Contact Email
Contact Phone
Last Modified Date
Description
Info URL

▼ API (Enable OAuth Settings)

Consumer Key and Secret [Manage Consumer Details](#)

Selected OAuth Scopes Full access (full)

Callback URL <https://login.salesforce.com/services/oauth2/token>

Enable for Device Flow

Require Proof Key for Code Exchange (PKCE) Extension for Supported Authorization Flows

8. On the **Connected App Name** page, do the following:

- From **Consumer Details**, copy and save the following in a text editor of your choice:
 - **Consumer Key** – You will need this to connect Salesforce Online to Amazon Q.
 - **Consumer Secret** – You will need this to connect Salesforce Online to Amazon Q.
 - Select **Apply**.

Connected App Name

[« Back to Manage Connected Apps](#)

Consumer Details

1 Consumer Key

2 Consumer Secret

Staged Consumer Details

Generate staged values for the consumer key and secret. When you apply the staged values, they replace the original consumer details.

Staged Consumer Key	Not generated
Staged Consumer Secret	Not generated

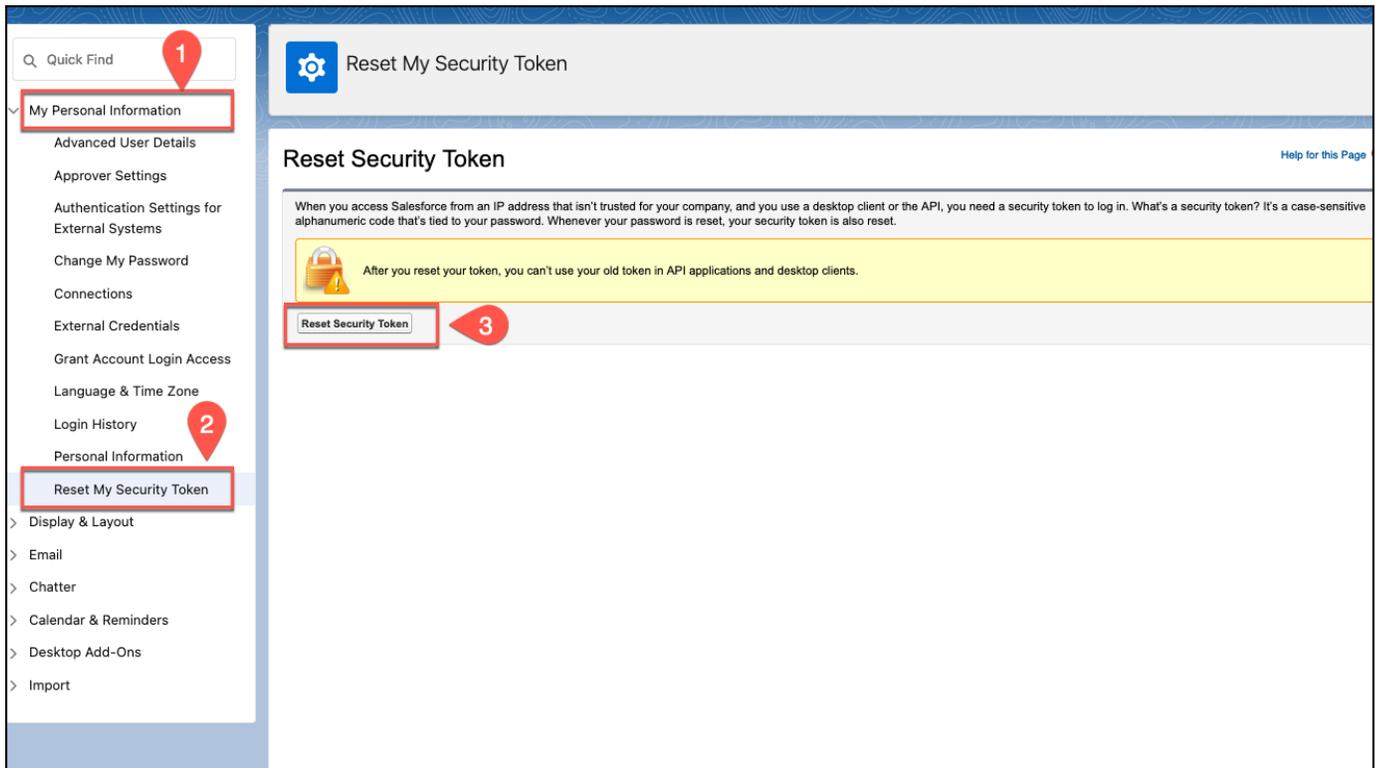
3

- Next, you will generate a security token. Navigate back to your Salesforce Online account home page. From the Salesforce Online profile menu, select **Settings**.

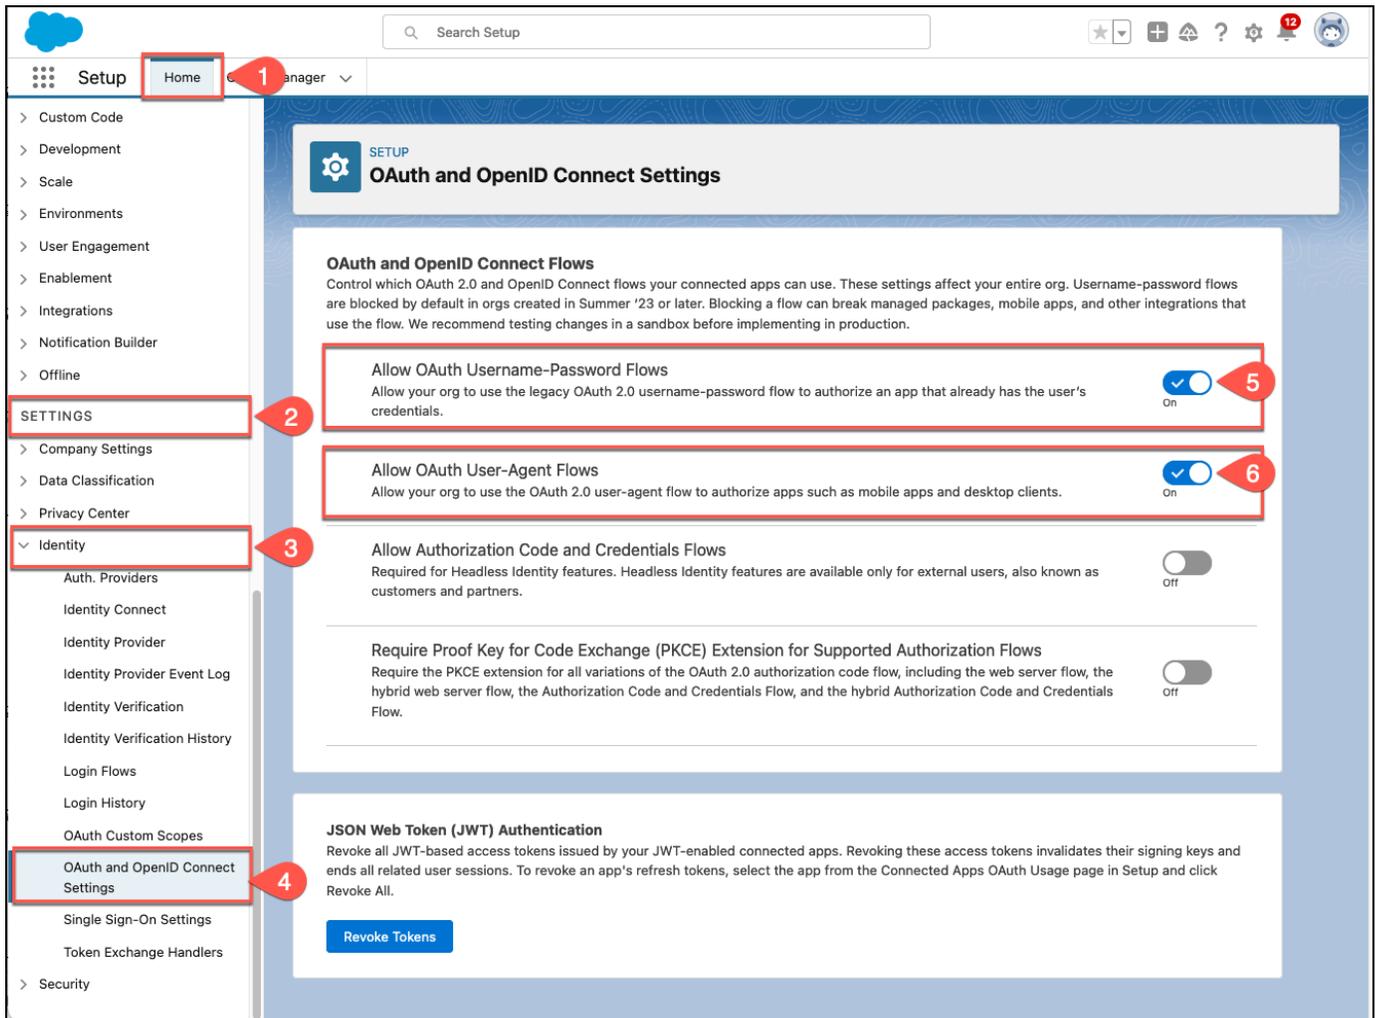
The screenshot shows the Salesforce Setup Home page. The top navigation bar includes a search bar for 'Search Setup' and a user profile icon. The user menu is open, showing options like 'Settings' and 'Log Out'. The main content area features three cards: 'Get Started with Einstein Bots', 'Mobile Publisher', and 'Real-time Collaborative Docs'. A 'Most Recently Used' section is visible at the bottom, showing a table with columns for NAME, TYPE, and OBJECT.

NAME	TYPE	OBJECT

- Then, from the left navigation menu, select **My Personal Information**. Then, select **Reset My Security Token**. Your security token will be sent to the email address connect to your Salesforce Online instance. You need this security token to connect Salesforce Online to Amazon Q.



11. Then, you activate OAuth Username-Password Flow for the Salesforce Online Connected App you've created. From the left navigation menu, from **Settings**, select **Identity** and then select **OAuth and OpenID Connect Settings**.
12. On the **OAuth and OpenID Connect Settings**, in **OAuth and OpenID Connect Flows**, make sure that both **Allow OAuth Username-Password Flows** and **Allow OAuth User-Agent Flows** are activated.



You now have the Salesforce Online host URL, username, password, security token, client ID, client secret, and authentication URL you need to connect Salesforce Online to Amazon Q.

Connecting Amazon Q Business to Salesforce Online using the console

The following procedure outlines how to connect Amazon Q Business to Salesforce Online using the AWS Management Console.

Connecting Amazon Q to Salesforce Online

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).

4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **Salesforce Online** page, enter the following information:

6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. In **Source**, enter the following information:

- **Salesforce URL** – Enter your Salesforce server URL. For example, *https://mysite.salesforce.com*.

8. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

 **Note**

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

9. **Authentication** – Enter the following information for your **AWS Secrets Manager secret**.
 - a. **Secret name** – A name for your secret.
 - b. For **Username**, **Password**, **Security token**, **Consumer key**, **Consumer secret**, and **Authentication URL** – Enter the authentication credential values that you created in your Salesforce account.
 - c. Choose **Save and add secret**.
10. **Configure VPC and security group** – *optional* – Choose whether you want to use a VPC. If you do, enter the following information:
 - a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
 - b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

11. **Identity crawler** – Choose to activate Amazon Q identity crawler to sync identity information. For more information, see [Identity crawler](#).
12. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

13. **Sync scope** – Set the content that you want to sync.
 - a. For **Standard objects**, **Standard objects with attachments**, **Standard objects without attachments**, and **Knowledge articles** – Select Salesforce entities or content types you want to crawl.

 **Note**

You must provide configuration information for indexing at least one of standard objects, knowledge articles, or chatter feeds. If you choose to crawl **Knowledge articles** you must specify the types of knowledge articles to index, the name of the articles, and whether to index the standard fields of all knowledge articles or only the fields of a custom article type. If you choose to index custom articles, you must specify the internal name of the article type. You can specify up to 10 article types.

- b. For **Custom objects** – Add custom object names. You can choose to include custom object attachments as well.
14. In **Additional configuration – optional**:
 - For **Entity regex patterns** and **Attachment regex patterns** – Add regular expression patterns to include or exclude certain files. You can add up to 100 patterns.
15. In **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.
 - **Full sync** – Sync all content regardless of the previous sync status.
 - **New or modified content sync** – Sync only new and modified documents.
 - **New, modified, or deleted content sync** – Sync only new, modified, and deleted documents.

For more details, see [Sync mode](#).

16. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
17. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
18. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:
 - a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
 - b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

 **Note**

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

19. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

20. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

 **Note**

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to Salesforce using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the configuration parameter to provide a JSON schema with all other configuration information specific to your data source connector.

For an example of the API request, see [CreateDataSource](#) in the Amazon Q API Reference.

Salesforce JSON schema

The following is the Salesforce JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    {
      "connectionConfiguration": {
        "type": "object",
        "properties": {
          {
            "repositoryEndpointMetadata": {
              {
                "type": "object",
                "properties": {
                  {
                    "hostUrl": {
                      {
                        "type": "string",
                        "pattern": "^https:\\\\[a-zA-Z0-9-\\.]*\\.\\.(salesforce|force).com\\/?$"
                      }
                    },
                    },
                  },
                "required": [
                  "hostUrl"
                ]
              }
            },
            },
          "required": [
            "repositoryEndpointMetadata"
          ]
        },
      }
    }
  }
}
```

```
"repositoryConfigurations": {
  "type": "object",
  "properties":
  {
    "account":
    {
      "type": "object",
      "properties":
      {
        "fieldMappings":
        {
          "type": "array",
          "items":
          [
            {
              "type": "object",
              "properties":
              {
                "indexFieldName":
                {
                  "type": "string"
                },
                "indexFieldType":
                {
                  "type": "string",
                  "enum":
                  [
                    "STRING",
                    "STRING_LIST",
                    "DATE",
                    "LONG"
                  ]
                },
                "dataSourceFieldName":
                {
                  "type": "string"
                },
                "dateFieldFormat":
                {
                  "type": "string",
                  "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
                }
              }
            }
          ]
        },
        "required":
```

```
        [
          "indexFieldName",
          "indexFieldType",
          "dataSourceFieldName"
        ]
      }
    ]
  },
  "required":
  [
    "fieldMappings"
  ],
  "contact":
  {
    "type": "object",
    "properties":
    {
      "fieldMappings":
      {
        "type": "array",
        "items":
        [
          {
            "type": "object",
            "properties":
            {
              "indexFieldName":
              {
                "type": "string"
              },
              "indexFieldType":
              {
                "type": "string",
                "enum":
                [
                  "STRING",
                  "STRING_LIST",
                  "DATE"
                ]
              },
              "dataSourceFieldName":
              {
```

```

        "type": "string"
      },
      "dateFieldFormat":
      {
        "type": "string",
        "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
      }
    },
    "required":
    [
      "indexFieldName",
      "indexFieldType",
      "dataSourceFieldName"
    ]
  }
]
}
},
"required":
[
  "fieldMappings"
]
},
"campaign":
{
  "type": "object",
  "properties":
  {
    "fieldMappings":
    {
      "type": "array",
      "items":
      [
        {
          "type": "object",
          "properties":
          {
            "indexFieldName":
            {
              "type": "string"
            },
            "indexFieldType":
            {
              "type": "string",

```

```
        "enum":
        [
            "STRING",
            "STRING_LIST",
            "DATE",
            "LONG"
        ]
    },
    "dataSourceFieldName":
    {
        "type": "string"
    },
    "dateFieldFormat":
    {
        "type": "string",
        "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
    }
},
"required":
[
    "indexFieldName",
    "indexFieldType",
    "dataSourceFieldName"
]
}
]
}
},
"required":
[
    "fieldMappings"
]
},
"case":
{
    "type": "object",
    "properties":
    {
        "fieldMappings":
        {
            "type": "array",
            "items":
            [
                {
```

```
    "type": "object",
    "properties":
    {
      "indexFieldName":
      {
        "type": "string"
      },
      "indexFieldType":
      {
        "type": "string",
        "enum":
        [
          "STRING",
          "STRING_LIST",
          "DATE"
        ]
      },
      "dataSourceFieldName":
      {
        "type": "string"
      },
      "dateFieldFormat":
      {
        "type": "string",
        "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
      }
    },
    "required":
    [
      "indexFieldName",
      "indexFieldType",
      "dataSourceFieldName"
    ]
  }
]
},
"required":
[
  "fieldMappings"
]
},
"product":
{
```

```
"type": "object",
"properties":
{
  "fieldMappings":
  {
    "type": "array",
    "items":
    [
      {
        "type": "object",
        "properties":
        {
          "indexFieldName":
          {
            "type": "string"
          },
          "indexFieldType":
          {
            "type": "string",
            "enum":
            [
              "STRING",
              "STRING_LIST",
              "DATE"
            ]
          },
          "dataSourceFieldName":
          {
            "type": "string"
          },
          "dateFieldFormat":
          {
            "type": "string",
            "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
          }
        }
      },
      "required":
      [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
      ]
    ]
  }
}
```

```
    }
  },
  "required":
  [
    "fieldMappings"
  ]
},
"lead":
{
  "type": "object",
  "properties":
  {
    "fieldMappings":
    {
      "type": "array",
      "items":
      [
        {
          "type": "object",
          "properties":
          {
            "indexFieldName":
            {
              "type": "string"
            },
            "indexFieldType":
            {
              "type": "string",
              "enum":
              [
                "STRING",
                "STRING_LIST",
                "DATE",
                "LONG"
              ]
            },
            "dataSourceFieldName":
            {
              "type": "string"
            },
            "dateFieldFormat":
            {
              "type": "string",
              "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
            }
          }
        }
      ]
    }
  }
}
```



```
        }
      },
      "required":
      [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
      ]
    }
  ]
}
},
"required":
[
  "fieldMappings"
],
"contract":
{
  "type": "object",
  "properties":
  {
    "fieldMappings":
    {
      "type": "array",
      "items":
      [
        {
          "type": "object",
          "properties":
          {
            "indexFieldName":
            {
              "type": "string"
            },
            "indexFieldType":
            {
              "type": "string",
              "enum":
              [
                "STRING",
                "STRING_LIST",
                "DATE"
              ]
            }
          }
        }
      ]
    }
  }
}
```

```
    },
    "dataSourceFieldName":
    {
      "type": "string"
    },
    "dateFieldFormat":
    {
      "type": "string",
      "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
    }
  },
  "required":
  [
    "indexFieldName",
    "indexFieldType",
    "dataSourceFieldName"
  ]
}
]
}
},
"required":
[
  "fieldMappings"
]
},
"partner":
{
  "type": "object",
  "properties":
  {
    "fieldMappings":
    {
      "type": "array",
      "items":
      [
        {
          "type": "object",
          "properties":
          {
            "indexFieldName":
            {
              "type": "string"
            }
          }
        }
      ]
    }
  }
}
```

```

        "indexFieldType":
        {
            "type": "string",
            "enum":
            [
                "STRING",
                "STRING_LIST",
                "DATE"
            ]
        },
        "dataSourceFieldName":
        {
            "type": "string"
        },
        "dateFieldFormat":
        {
            "type": "string",
            "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
        }
    },
    "required":
    [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
    ]
}
]
}
},
"required":
[
    "fieldMappings"
]
},
"profile":
{
    "type": "object",
    "properties":
    {
        "fieldMappings":
        {
            "type": "array",
            "items":

```

```
[
  {
    "type": "object",
    "properties":
    {
      "indexFieldName":
      {
        "type": "string"
      },
      "indexFieldType":
      {
        "type": "string",
        "enum":
        [
          "STRING",
          "STRING_LIST",
          "DATE"
        ]
      },
      "dataSourceFieldName":
      {
        "type": "string"
      },
      "dateFieldFormat":
      {
        "type": "string",
        "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
      }
    },
    "required":
    [
      "indexFieldName",
      "indexFieldType",
      "dataSourceFieldName"
    ]
  }
],
"required":
[
  "fieldMappings"
],
},
```

```
"idea":
{
  "type": "object",
  "properties":
  {
    "fieldMappings":
    {
      "type": "array",
      "items":
      [
        {
          "type": "object",
          "properties":
          {
            "indexFieldName":
            {
              "type": "string"
            },
            "indexFieldType":
            {
              "type": "string",
              "enum":
              [
                "STRING",
                "STRING_LIST",
                "DATE",
                "LONG"
              ]
            },
            "dataSourceFieldName":
            {
              "type": "string"
            },
            "dateFieldFormat":
            {
              "type": "string",
              "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
            }
          }
        },
        "required":
        [
          "indexFieldName",
          "indexFieldType",
          "dataSourceFieldName"
        ]
      ]
    }
  }
}
```

```
        ]
      }
    ]
  },
  "required":
  [
    "fieldMappings"
  ]
},
"pricebook":
{
  "type": "object",
  "properties":
  {
    "fieldMappings":
    {
      "type": "array",
      "items":
      [
        {
          "type": "object",
          "properties":
          {
            "indexFieldName":
            {
              "type": "string"
            },
            "indexFieldType":
            {
              "type": "string",
              "enum":
              [
                "STRING",
                "STRING_LIST",
                "DATE"
              ]
            },
            "dataSourceFieldName":
            {
              "type": "string"
            },
            "dateFieldFormat":
            {
```

```

        "type": "string",
        "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
    }
},
"required":
[
    "indexFieldName",
    "indexFieldType",
    "dataSourceFieldName"
]
}
]
}
},
"required":
[
    "fieldMappings"
]
},
"task":
{
    "type": "object",
    "properties":
    {
        "fieldMappings":
        {
            "type": "array",
            "items":
            [
                {
                    "type": "object",
                    "properties":
                    {
                        "indexFieldName":
                        {
                            "type": "string"
                        },
                        "indexFieldType":
                        {
                            "type": "string",
                            "enum":
                            [
                                "STRING",
                                "STRING_LIST",

```

```
        "DATE"
      ]
    },
    "dataSourceFieldName":
    {
      "type": "string"
    },
    "dateFieldFormat":
    {
      "type": "string",
      "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
    }
  },
  "required":
  [
    "indexFieldName",
    "indexFieldType",
    "dataSourceFieldName"
  ]
}
]
}
},
"required":
[
  "fieldMappings"
]
},
"solution":
{
  "type": "object",
  "properties":
  {
    "fieldMappings":
    {
      "type": "array",
      "items":
      [
        {
          "type": "object",
          "properties":
          {
            "indexFieldName":
            {
```



```

        "type": "string"
      },
      "indexFieldType":
      {
        "type": "string",
        "enum":
        [
          "STRING",
          "STRING_LIST",
          "DATE"
        ]
      },
      "dataSourceFieldName":
      {
        "type": "string"
      },
      "dateFieldFormat":
      {
        "type": "string",
        "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
      }
    },
    "required":
    [
      "indexFieldName",
      "indexFieldType",
      "dataSourceFieldName"
    ]
  }
]
}
},
"required":
[
  "fieldMappings"
]
},
"attachment":
{
  "type": "object",
  "properties":
  {
    "fieldMappings":
    {

```

```
    "type": "array",
    "items":
    [
      {
        "type": "object",
        "properties":
        {
          "indexFieldName":
          {
            "type": "string"
          },
          "indexFieldType":
          {
            "type": "string",
            "enum":
            [
              "STRING",
              "STRING_LIST",
              "DATE",
              "LONG"
            ]
          },
          "dataSourceFieldName":
          {
            "type": "string"
          },
          "dateFieldFormat":
          {
            "type": "string",
            "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
          }
        },
        "required":
        [
          "indexFieldName",
          "indexFieldType",
          "dataSourceFieldName"
        ]
      }
    ]
  },
  "required":
  [
```

```
    "fieldMappings"
  ]
},
"user":
{
  "type": "object",
  "properties":
  {
    "fieldMappings":
    {
      "type": "array",
      "items":
      [
        {
          "type": "object",
          "properties":
          {
            "indexFieldName":
            {
              "type": "string"
            },
            "indexFieldType":
            {
              "type": "string",
              "enum":
              [
                "STRING",
                "STRING_LIST",
                "DATE"
              ]
            },
            "dataSourceFieldName":
            {
              "type": "string"
            },
            "dateFieldFormat":
            {
              "type": "string",
              "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
            }
          }
        },
        "required":
        [
          "indexFieldName",
```

```
        "indexFieldType",
        "dataSourceFieldName"
      ]
    }
  ]
},
"required":
[
  "fieldMappings"
],
"document":
{
  "type": "object",
  "properties":
  {
    "fieldMappings":
    {
      "type": "array",
      "items":
      [
        {
          "type": "object",
          "properties":
          {
            "indexFieldName":
            {
              "type": "string"
            },
            "indexFieldType":
            {
              "type": "string",
              "enum":
              [
                "STRING",
                "STRING_LIST",
                "DATE",
                "LONG"
              ]
            },
            "dataSourceFieldName":
            {
              "type": "string"
            }
          }
        }
      ]
    }
  }
}
```

```
    },
    "dateFieldFormat":
    {
      "type": "string",
      "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
    }
  },
  "required":
  [
    "indexFieldName",
    "indexFieldType",
    "dataSourceFieldName"
  ]
}
]
}
},
"required":
[
  "fieldMappings"
]
},
"knowledgeArticles":
{
  "type": "object",
  "properties":
  {
    "fieldMappings":
    {
      "type": "array",
      "items":
      [
        {
          "type": "object",
          "properties":
          {
            "indexFieldName":
            {
              "type": "string"
            },
            "indexFieldType":
            {
              "type": "string",
              "enum":
```

```

        [
            "STRING",
            "STRING_LIST",
            "DATE"
        ]
    },
    "dataSourceFieldName":
    {
        "type": "string"
    },
    "dateFieldFormat":
    {
        "type": "string",
        "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
    }
},
"required":
[
    "indexFieldName",
    "indexFieldType",
    "dataSourceFieldName"
]
}
]
}
},
"required":
[
    "fieldMappings"
]
},
"group":
{
    "type": "object",
    "properties":
    {
        "fieldMappings":
        {
            "type": "array",
            "items":
            [
                {
                    "type": "object",
                    "properties":

```

```
    {
      "indexFieldName":
      {
        "type": "string"
      },
      "indexFieldType":
      {
        "type": "string",
        "enum":
        [
          "STRING",
          "STRING_LIST",
          "DATE"
        ]
      },
      "dataSourceFieldName":
      {
        "type": "string"
      },
      "dateFieldFormat":
      {
        "type": "string",
        "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
      }
    },
    "required":
    [
      "indexFieldName",
      "indexFieldType",
      "dataSourceFieldName"
    ]
  }
]
},
"required":
[
  "fieldMappings"
]
},
"opportunity":
{
  "type": "object",
  "properties":
```

```
{
  "fieldMappings":
  {
    "type": "array",
    "items":
    [
      {
        "type": "object",
        "properties":
        {
          "indexFieldName":
          {
            "type": "string"
          },
          "indexFieldType":
          {
            "type": "string",
            "enum":
            [
              "STRING",
              "STRING_LIST",
              "DATE",
              "LONG"
            ]
          },
          "dataSourceFieldName":
          {
            "type": "string"
          },
          "dateFieldFormat":
          {
            "type": "string",
            "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
          }
        },
        "required":
        [
          "indexFieldName",
          "indexFieldType",
          "dataSourceFieldName"
        ]
      }
    ]
  }
}
```



```
    },
    "required":
    [
      "fieldMappings"
    ]
  },
  "chatter":
  {
    "type": "object",
    "properties":
    {
      "fieldMappings":
      {
        "type": "array",
        "items":
        [
          {
            "type": "object",
            "properties":
            {
              "indexFieldName":
              {
                "type": "string"
              },
              "indexFieldType":
              {
                "type": "string",
                "enum":
                [
                  "STRING",
                  "STRING_LIST",
                  "DATE"
                ]
              },
              "dataSourceFieldName":
              {
                "type": "string"
              },
              "dateFieldFormat":
              {
                "type": "string",
                "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
              }
            }
          }
        ]
      }
    }
  },
```

```
        "required":
          [
            "indexFieldName",
            "indexFieldType",
            "dataSourceFieldName"
          ]
        }
      ]
    },
    "required":
    [
      "fieldMappings"
    ]
  },
  "customEntity":
  {
    "type": "object",
    "properties":
    {
      "fieldMappings":
      {
        "type": "array",
        "items":
        [
          {
            "type": "object",
            "properties":
            {
              "indexFieldName":
              {
                "type": "string"
              },
              "indexFieldType":
              {
                "type": "string",
                "enum":
                [
                  "STRING",
                  "STRING_LIST",
                  "DATE"
                ]
              },
            }
          }
        ]
      },
      "dataSourceFieldName":
```

```

        {
            "type": "string"
        },
        "dateFieldFormat":
        {
            "type": "string",
            "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
        }
    },
    "required":
    [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
    ]
}
]
}
},
"required":
[
    "fieldMappings"
]
}
}
},
"additionalProperties": {
    "type": "object",
    "properties":
    {
        "accountFilter":{
            "type": "array",
            "items":
            {
                "type": "string"
            }
        },
        "contactFilter":{
            "type": "array",
            "items":
            {
                "type": "string"
            }
        }
    },
}
},

```

```
"caseFilter":{
  "type": "array",
  "items":
  {
    "type": "string"
  }
},
"campaignFilter":{
  "type": "array",
  "items":
  {
    "type": "string"
  }
},
"contractFilter":{
  "type": "array",
  "items":
  {
    "type": "string"
  }
},
"groupFilter":{
  "type": "array",
  "items":
  {
    "type": "string"
  }
},
"leadFilter":{
  "type": "array",
  "items":
  {
    "type": "string"
  }
},
"productFilter":{
  "type": "array",
  "items":
  {
    "type": "string"
  }
},
"opportunityFilter":{
  "type": "array",
```

```
    "items":
      {
        "type": "string"
      }
  },
  "partnerFilter":{
    "type": "array",
    "items":
      {
        "type": "string"
      }
  },
  "pricebookFilter":{
    "type": "array",
    "items":
      {
        "type": "string"
      }
  },
  "ideaFilter":{
    "type": "array",
    "items":
      {
        "type": "string"
      }
  },
  "profileFilter":{
    "type": "array",
    "items":
      {
        "type": "string"
      }
  },
  "taskFilter":{
    "type": "array",
    "items":
      {
        "type": "string"
      }
  },
  "solutionFilter":{
    "type": "array",
    "items":
      {
```

```
    "type": "string"
  }
},
"userFilter":{
  "type": "array",
  "items":
  {
    "type": "string"
  }
},
"chatterFilter":{
  "type": "array",
  "items":
  {
    "type": "string"
  }
},
"documentFilter":{
  "type": "array",
  "items":
  {
    "type": "string"
  }
},
"knowledgeArticleFilter":{
  "type": "array",
  "items":
  {
    "type": "string"
  }
},
"customEntities":{
  "type": "array",
  "items":
  {
    "type": "string"
  }
},
"isCrawlAcl": {
  "type": "boolean"
},
"fieldForUserId": {
  "type": "string"
},
}
```

```
"isCrawlAccount": {
  "type": "boolean"
},
"isCrawlContact": {
  "type": "boolean"
},
"isCrawlCase": {
  "type": "boolean"
},
"isCrawlCampaign": {
  "type": "boolean"
},
"isCrawlProduct": {
  "type": "boolean"
},
"isCrawlLead": {
  "type": "boolean"
},
"isCrawlContract": {
  "type": "boolean"
},
"isCrawlPartner": {
  "type": "boolean"
},
"isCrawlProfile": {
  "type": "boolean"
},
"isCrawlIdea": {
  "type": "boolean"
},
"isCrawlPricebook": {
  "type": "boolean"
},
"isCrawlDocument": {
  "type": "boolean"
},
"crawlSharedDocument": {
  "type": "boolean"
},
"isCrawlGroup": {
  "type": "boolean"
},
"isCrawlOpportunity": {
  "type": "boolean"
}
```

```
    },
    "isCrawlChatter": {
      "type": "boolean"
    },
    "isCrawlUser": {
      "type": "boolean"
    },
    "isCrawlSolution":{
      "type": "boolean"
    },
    "isCrawlTask":{
      "type": "boolean"
    },
    },

    "isCrawlAccountAttachments": {
      "type": "boolean"
    },
    },
    "isCrawlContactAttachments": {
      "type": "boolean"
    },
    },
    "isCrawlCaseAttachments": {
      "type": "boolean"
    },
    },
    "isCrawlCampaignAttachments": {
      "type": "boolean"
    },
    },
    "isCrawlLeadAttachments": {
      "type": "boolean"
    },
    },
    "isCrawlContractAttachments": {
      "type": "boolean"
    },
    },
    "isCrawlGroupAttachments": {
      "type": "boolean"
    },
    },
    "isCrawlOppportunityAttachments": {
      "type": "boolean"
    },
    },
    "isCrawlChatterAttachments": {
      "type": "boolean"
    },
    },
    "isCrawlSolutionAttachments":{
      "type": "boolean"
    },
    },
```



```
"isCrawlTaskAttachments":{
  "type": "boolean"
},
"isCrawlCustomEntityAttachments":{
  "type": "boolean"
},
"isCrawlKnowledgeArticles": {
  "type": "object",
  "properties":
  {
    "isCrawlDraft": {
      "type": "boolean"
    },
    "isCrawlPublish": {
      "type": "boolean"
    },
    "isCrawlArchived": {
      "type": "boolean"
    }
  }
},
"inclusionDocumentFileTypePatterns":{
  "type": "array",
  "items":
  {
    "type": "string"
  }
},
"exclusionDocumentFileTypePatterns": {
  "type": "array",
  "items":
  {
    "type": "string"
  }
},
"inclusionDocumentFileNamePatterns":{
  "type": "array",
  "items":
  {
    "type": "string"
  }
},
"exclusionDocumentFileNamePatterns": {
  "type": "array",
```

```
    "items":
      {
        "type": "string"
      }
  },
  "inclusionAccountFileTypePatterns": {
    "type": "array",
    "items":
      {
        "type": "string"
      }
  },
  "exclusionAccountFileTypePatterns": {
    "type": "array",
    "items":
      {
        "type": "string"
      }
  },
  "inclusionAccountFileNamePatterns":{
    "type": "array",
    "items":
      {
        "type": "string"
      }
  },
  "exclusionAccountFileNamePatterns":{
    "type": "array",
    "items":
      {
        "type": "string"
      }
  },
  "inclusionCampaignFileTypePatterns": {
    "type": "array",
    "items":
      {
        "type": "string"
      }
  },
  "exclusionCampaignFileTypePatterns": {
    "type": "array",
    "items":
      {
```

```
    "type": "string"
  }
},
"inclusionCampaignFileNamePatterns":{
  "type": "array",
  "items":
  {
    "type": "string"
  }
},
"exclusionCampaignFileNamePatterns":{
  "type": "array",
  "items":
  {
    "type": "string"
  }
},
"inclusionCaseFileTypePatterns":{
  "type": "array",
  "items":
  {
    "type": "string"
  }
},
"exclusionCaseFileTypePatterns":{
  "type": "array",
  "items":
  {
    "type": "string"
  }
},
"inclusionCaseFileNamePatterns":{
  "type": "array",
  "items":
  {
    "type": "string"
  }
},
"exclusionCaseFileNamePatterns":{
  "type": "array",
  "items":
  {
    "type": "string"
  }
}
```

```
  },
  "inclusionContactFileTypePatterns":{
    "type": "array",
    "items":
    {
      "type": "string"
    }
  },
  "exclusionContactFileTypePatterns":{
    "type": "array",
    "items":
    {
      "type": "string"
    }
  },
  "inclusionContactFileNamePatterns":{
    "type": "array",
    "items":
    {
      "type": "string"
    }
  },
  "exclusionContactFileNamePatterns":{
    "type": "array",
    "items":
    {
      "type": "string"
    }
  },
  "inclusionContractFileTypePatterns":{
    "type": "array",
    "items":
    {
      "type": "string"
    }
  },
  "exclusionContractFileTypePatterns":{
    "type": "array",
    "items":
    {
      "type": "string"
    }
  },
  "inclusionContractFileNamePatterns":{
```

```
    "type": "array",
    "items":
    {
      "type": "string"
    }
  },
  "exclusionContractFileNamePatterns":{
    "type": "array",
    "items":
    {
      "type": "string"
    }
  },
  "inclusionLeadFileTypePatterns":{
    "type": "array",
    "items":
    {
      "type": "string"
    }
  },
  "exclusionLeadFileTypePatterns":{
    "type": "array",
    "items":
    {
      "type": "string"
    }
  },
  "inclusionLeadFileNamePatterns":{
    "type": "array",
    "items":
    {
      "type": "string"
    }
  },
  "exclusionLeadFileNamePatterns":{
    "type": "array",
    "items":
    {
      "type": "string"
    }
  },
  "inclusionOppportunityFileTypePatterns":{
    "type": "array",
    "items":
```

```
{
  "type": "string"
},
"exclusionOpportunityFileTypePatterns":{
  "type": "array",
  "items":
  {
    "type": "string"
  }
},
"inclusionOpportunityFileNamePatterns":{
  "type": "array",
  "items":
  {
    "type": "string"
  }
},
"exclusionOpportunityFileNamePatterns":{
  "type": "array",
  "items":
  {
    "type": "string"
  }
},
"inclusionSolutionFileTypePatterns":{
  "type": "array",
  "items":
  {
    "type": "string"
  }
},
"exclusionSolutionFileTypePatterns":{
  "type": "array",
  "items":
  {
    "type": "string"
  }
},
"inclusionSolutionFileNamePatterns":{
  "type": "array",
  "items":
  {
    "type": "string"
  }
}
```

```
    }
  },
  "exclusionSolutionFileNamePatterns":{
    "type": "array",
    "items":
    {
      "type": "string"
    }
  },
  "inclusionTaskFileTypePatterns":{
    "type": "array",
    "items":
    {
      "type": "string"
    }
  },
  "exclusionTaskFileTypePatterns":{
    "type": "array",
    "items":
    {
      "type": "string"
    }
  },
  "inclusionTaskFileNamePatterns":{
    "type": "array",
    "items":
    {
      "type": "string"
    }
  },
  "exclusionTaskFileNamePatterns":{
    "type": "array",
    "items":
    {
      "type": "string"
    }
  },
  "inclusionGroupFileTypePatterns":{
    "type": "array",
    "items":
    {
      "type": "string"
    }
  },
},
```

```
"exclusionGroupFileTypePatterns":{
  "type": "array",
  "items":
  {
    "type": "string"
  }
},
"inclusionGroupFileNamePatterns":{
  "type": "array",
  "items":
  {
    "type": "string"
  }
},
"exclusionGroupFileNamePatterns":{
  "type": "array",
  "items":
  {
    "type": "string"
  }
},
"inclusionChatterFileTypePatterns":{
  "type": "array",
  "items":
  {
    "type": "string"
  }
},
"exclusionChatterFileTypePatterns":{
  "type": "array",
  "items":
  {
    "type": "string"
  }
},
"inclusionChatterFileNamePatterns":{
  "type": "array",
  "items":
  {
    "type": "string"
  }
},
"exclusionChatterFileNamePatterns":{
  "type": "array",
```



```
    "items":
      {
        "type": "string"
      }
  },
  "inclusionCustomEntityTypePatterns":{
    "type": "array",
    "items":
      {
        "type": "string"
      }
  },
  "exclusionCustomEntityTypePatterns":{
    "type": "array",
    "items":
      {
        "type": "string"
      }
  },
  "inclusionCustomEntityFileNamePatterns":{
    "type": "array",
    "items":
      {
        "type": "string"
      }
  },
  "exclusionCustomEntityFileNamePatterns":{
    "type": "array",
    "items":
      {
        "type": "string"
      }
  }
},
"required":
[]
},
"enableIdentityCrawler": {
  "type": "boolean"
},
"type": {
  "type": "string",
"enum": [
  "SALESFORCEV2",
```

```

    "SALESFORCE"
  ]
},
"syncMode": {
  "type": "string",
  "enum": [
    "FULL_CRAWL",
    "FORCED_FULL_CRAWL",
    "CHANGE_LOG"
  ]
},
"secretArn": {
  "type": "string",
  "minLength": 20,
  "maxLength": 2048
}
},
"version": {
  "type": "string",
  "anyOf": [
    {
      "pattern": "1.0.0"
    }
  ]
},
"required": [
  "connectionConfiguration",
  "repositoryConfigurations",
  "syncMode",
  "additionalProperties",
  "secretArn",
  "type"
]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.

Configuration	Description
repositoryEndpointMetadata	The endpoint information for the data source.
hostUrl	The URL of the Salesforce instance to be indexed.
repositoryConfigurations	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings.
<ul style="list-style-type: none"> • account • contact • campaign • case • product • lead • contract • partner • profile • idea • pricebook • task • solution • attachment • user • document • knowledgeArticles • group • opportunity • chatter • customEntity 	A list of objects that map the attributes or field names of your Salesforce entities to Amazon Q index field names.

Configuration	Description
additionalProperties	Additional configuration options for your content in your data source.
<ul style="list-style-type: none"> • accountFilter • contactFilter • caseFilter • campaignFilter • contractFilter • groupFilter • leadFilter • productFilter • opportunityFilter • partnerFilter • pricebookFilter • ideaFilter • profileFilter • taskFilter • solutionFilter • userFilter • chatterFilter • documentFilter • knowledgeArticleFilter 	Filters to specify content for Amazon Q to crawl.
customEntities	Custom entities that Amazon Q should crawl.

Configuration	Description
<p><code>inclusionPatterns</code></p> <ul style="list-style-type: none"> • <code>inclusionDocumentFileTypePatterns</code> • <code>inclusionDocumentFileNamePatterns</code> • <code>inclusionAccountFileTypePatterns</code> • <code>inclusionCampaignFileTypePatterns</code> • <code>inclusionDocumentFileNamePatterns</code> • <code>inclusionCampaignFileNamePatterns</code> • <code>inclusionCaseFileTypePatterns</code> • <code>inclusionCaseFileNamePatterns</code> • <code>inclusionContactFileTypePatterns</code> • <code>inclusionContractFileNamePatterns</code> • <code>inclusionLeadFileTypePatterns</code> • <code>inclusionLeadFileNamePatterns</code> • <code>inclusionOpportunityFileTypePatterns</code> • <code>inclusionOpportunityFileNamePatterns</code> • <code>inclusionSolutionFileTypePatterns</code> • <code>inclusionSolutionFileNamePatterns</code> • <code>inclusionTaskFileTypePatterns</code> 	<p>A list of regular expression patterns to <i>include</i> specific files in your Salesforce data source. Files that match the patterns are included in the index. Files that don't match the patterns are excluded from the index. If a file matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence and the file isn't included in the index.</p>

Configuration	Description
<ul style="list-style-type: none">• inclusionTaskFileNamePatterns• inclusionGroupFileTypePatterns• inclusionGroupFileNamePatterns• inclusionChatterFileTypePatterns• inclusionChatterFileNamePatterns• inclusionCustomEntityTypePatterns• inclusionCustomEntityFileNamePatterns	

Configuration	Description
<p>exclusionPatterns</p> <ul style="list-style-type: none"> • exclusionDocumentFileTypePatterns • exclusionDocumentFileNamePatterns • exclusionAccountFileTypePatterns • exclusionCampaignFileTypePatterns • exclusionCampaignFileNamePatterns • exclusionCaseFileTypePatterns • exclusionCaseFileNamePatterns • exclusionContactFileTypePatterns • exclusionContractFileNamePatterns • exclusionLeadFileTypePatterns • exclusionLeadFileNamePatterns • exclusionOpportunityFileTypePatterns • exclusionOpportunityFileNamePatterns • exclusionSolutionFileTypePatterns • exclusionSolutionFileNamePatterns • exclusionTaskFileTypePatterns • exclusionTaskFileNamePatterns • exclusionGroupFileTypePatterns 	<p>A list of regular expression patterns to <i>exclude</i> specific files in your Salesforce data source. Files that match the patterns are excluded from the index. Files that don't match the patterns are included in the index. If a file matches both an exclusion and inclusion pattern, the exclusion pattern takes precedence and the file isn't included in the index.</p>

Configuration	Description
<ul style="list-style-type: none">• exclusionGroupFileNamePatterns• exclusionChatterFileTypePatterns• exclusionChatterFileNamePatterns• exclusionCustomEntityTypePatterns• exclusionCustomEntityFileNamePatterns	
isCrawlAcl	Specify true to crawl access control information from documents.
fieldForUserId	Specify field to use for UserId for ACL crawling.

Configuration	Description
<ul style="list-style-type: none"> • <code>isCrawlAccount</code> • <code>isCrawlContact</code> • <code>isCrawlCase</code> • <code>isCrawlCampaign</code> • <code>isCrawlProduct</code> • <code>isCrawlLead</code> • <code>isCrawlContract</code> • <code>isCrawlPartner</code> • <code>isCrawlProfile</code> • <code>isCrawlIdea</code> • <code>isCrawlPricebook</code> • <code>isCrawlDocument</code> • <code>crawlSharedDocument</code> • <code>isCrawlGroup</code> • <code>isCrawlOpportunity</code> • <code>isCrawlChatter</code> • <code>isCrawlUser</code> • <code>isCrawlSolution</code> • <code>isCrawlTask</code> • <code>isCrawlAccountAttachments</code> • <code>isCrawlContactAttachments</code> • <code>isCrawlCaseAttachments</code> • <code>isCrawlCampaignAttachments</code> • <code>isCrawlLeadAttachments</code> • <code>isCrawlContractAttachments</code> • <code>isCrawlGroupAttachments</code> • <code>isCrawlOpportunityAttachments</code> • <code>isCrawlChatterAttachments</code> 	<p>true to index corresponding files in your Salesforce account.</p>

Configuration	Description
<ul style="list-style-type: none"> • <code>isCrawlSolutionAttachments</code> • <code>isCrawlTaskAttachments</code> • <code>isCrawlCustomEntityAttachments</code> • <code>isCrawlKnowledgeArticles</code> <ul style="list-style-type: none"> • <code>isCrawlDraft</code> • <code>isCrawlPublish</code> • <code>isCrawlArchived</code> 	
<code>type</code>	The type of data source. Specify SALESFORCE as your data source type.
<code>enableIdentityCrawler</code>	true to activate identity crawler. Identity crawler is activated by default. Crawling identity information on users and groups with access to certain documents is useful for user context filtering. Search results are filtered based on the user or their group access to documents. See Identity crawler for more information.

Configuration	Description
syncMode	<p>Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose between the following options:</p> <ul style="list-style-type: none">• Use <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index• Use <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index• Use <code>CHANGE_LOG</code> to incrementally crawl only new and modified content each time your data source syncs with your index

Configuration	Description
secretARN	<p>The Amazon Resource Name (ARN) of an AWS Secrets Manager secret that contains the key-value pairs required to connect to your Salesforce data source. The secret must contain a JSON structure with the following keys:</p> <pre data-bbox="829 537 1507 1451"> { "authenticationUrl": "The OAUTH endpoint that Amazon Q connects to get an OAUTH token.", "consumerKey": "The application public key generated when you created your Salesforce application." , "consumerSecret": "The application private key generated when you created your Salesforce application." , "password": "The password associated with the user logging in to the Salesforce instance." , "securityToken": "The token associated with the user account logging in to the Salesforce instance." , "username": "The user name of the user logging in to the Salesforce instance." } </pre>
version	The version of this template that's currently supported.

How Amazon Q Business connector crawls Salesforce ACLs

When you connect an Salesforce data source to Amazon Q Business, Amazon Q Business crawls ACL information attached to a document (user and group information) from your Salesforce instance.

If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

You can apply ACL based chat filtering using Salesforce standard objects and chatter feeds. ACL based chat filtering isn't available for Salesforce knowledge articles.

For standard objects, the `_user_id` and `_group_ids` are used as follows:

- `_user_id` – The username of the Salesforce user.
- `_group_ids` – The group names in Salesforce.
 - Name of the Salesforce Profile
 - Name of the Salesforce Group
 - Name of the Salesforce UserRole
 - Name of the Salesforce PermissionSet

For chatter feeds, the `_user_id` and `_group_ids` are used as follows:

- `_user_id` – The username of the Salesforce user. Only available if the item is posted in the user's feed.
- `_group_ids` – Group IDs are used as follows. Only available if the feed item is posted in a chatter or collaboration group.
 - The name of the chatter or collaboration group.
 - If the group is public, PUBLIC:ALL.

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business Salesforce Online data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q Business enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a attribute mapping already available, or if you want to map additional document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q Salesforce connector supports the following entities and the associated reserved and custom attributes.

Note

You can map any Salesforce field to the document title or document body Amazon Q reserved/default index fields.

Supported entities and field mappings

- [Account](#)
- [Campaign](#)
- [Case](#)
- [Contact](#)
- [Contract](#)

- [Document](#)
- [Group](#)
- [Idea](#)
- [Lead](#)
- [Opportunity](#)
- [Partner](#)
- [Pricebook](#)
- [Product](#)
- [Solution](#)
- [Profile](#)
- [Task](#)
- [User](#)
- [Chatter](#)
- [Knowledge articles](#)
- [Attachments](#)
- [Custom object](#)

Account

Amazon Q supports crawling [Salesforce Online Accounts](#) and offers the following account field mappings.

Salesforce field name	Index field name	Description	Data type
category	_category	Default	String
sourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date
authors	_authors	Default	String list

Salesforce field name	Index field name	Description	Data type
lastModifiedBy	sf_last_modified_by	Custom	String
shippingCity	sf_shipping_city	Custom	String
shippingCountry	sf_shipping_country	Custom	String
shippingState	sf_shipping_state	Custom	String
website	sf_website	Custom	String
industry	sf_industry	Custom	String
accountSource	sf_account_source	Custom	String
billingCity	sf_billing_city	Custom	String
billingCountry	sf_billing_country	Custom	String
billingState	sf_billing_state	Custom	String
createdBy	sf_created_by	Custom	String
lastActivityDate	sf_last_activity_date	Custom	Date
parentId	sf_parent_id	Custom	String
typeValue	sf_type_value	Custom	String
billingStreet	sf_billing_street	Custom	String
billingPostalCode	sf_billing_postal_code	Custom	String
billingLatitude	sf_billing_latitude	Custom	String
billingLongitude	sf_billing_longitude	Custom	String
billingGeocodeAccuracy	sf_billing_geocode_accuracy	Custom	String

Salesforce field name	Index field name	Description	Data type
shippingStreet	sf_shipping_street	Custom	String
shippingPostalCode	sf_shipping_postal_code	Custom	String
phone	sf_phone	Custom	String
fax	sf_fax	Custom	String
annualRevenue	sf_annual_revenue	Custom	String
numberOfEmployees	sf_number_of_employees	Custom	Long (numeric)
jigsaw	sf_jigsaw	Custom	String

Campaign

Amazon Q supports crawling [Salesforce Online Campaigns](#) and offers the following campaign field mappings.

Salesforce field name	Index field name	Description	Data type
category	_category	Default	String
sourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
isActive	sf_is_active	Custom	String
updatedAt	_last_updated_at	Default	Date
ownerName	_authors	Default	String list
lastModifiedBy	sf_last_modified_by	Custom	String

Salesforce field name	Index field name	Description	Data type
createdBy	sf_created_by	Custom	String
lastActivityDate	sf_last_activity_date	Custom	Date
parentId	sf_parent_id	Custom	String
campaignName	sf_campaign_name	Custom	String
status	sf_status	Custom	String
parentName	sf_parent_name	Custom	String
campaignType	sf_type	Custom	String
expectedRevenue	sf_expected_revenue	Custom	Long (numeric)
budgetedCost	sf_budgeted_cost	Custom	Long (numeric)
actualCost	sf_actual_cost	Custom	Long(numeric)
expectedResponse	sf_expected_response	Custom	String
numberSent	sf_number_sent	Default	Long numeric)
numberOfLeads	sf_number_of_leads	Custom	Long (numeric)
numberOfConvertedLeads	sf_number_of_converted_leads	Custom	Long (numeric)
numberOfContacts	sf_number_of_contacts	Custom	Long (numeric)
numberOfResponses	sf_number_of_responses	Custom	Long (numeric)
numberOfOpportunities	sf_number_of_opportunities	Custom	Long (numeric)

Salesforce field name	Index field name	Description	Data type
numberOfWonOpportunities	sf_number_of_won_opportunities	Custom	Long (numeric)
amountAllOpportunities	sf_amount_all_opportunities	Custom	Long (numeric)
amountWonOpportunities	sf_amount_won_opportunities	Custom	Long (numeric)

Case

Amazon Q supports crawling [Salesforce Online Cases](#) and offers the following case field mappings.

Salesforce field name	Index field name	Description	Data type
category	_category	Default	String
sourceUrl	_source_uri	Default	String
authors	_authors	Default	String list
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date
ownerName	sf_owner_name	Custom	String
createdBy	sf_created_by	Custom	String
caseNumber	sf_case_number	Custom	String
isClosed	sf_is_closed	Custom	String
isEscalated	sf_is_escalated	Custom	String
priority	sf_priority	Custom	String

Salesforce field name	Index field name	Description	Data type
status	sf_status	Custom	String
accountName	sf_account_name	Custom	String
lastModifiedBy	af_last_modified_by	Custom	String
updatedAt	_last_updated_at	Default	Date
typeValue	sf_type	Custom	String
reason	sf_reason	Custom	String
contactId	sf_contact_id	Custom	String
origin	sf_origin	Custom	String
parentId	sf_parent_id	Custom	String
contactName	sf_contact_name	Custom	String
parentCaseNumber	sf_parent_case_number	Custom	String
parentSubject	sf_parent_subject	Custom	String
suppliedEmail	sf_supplied_email	Custom	String
contactPhone	sf_contact_phone	Custom	String
contactMobile	sf_contact_mobile	Custom	String
contactEmail	sf_contact_email	Custom	String
contactFax	sf_contact_fax	Custom	String
comments	sf_comments	Custom	String
lastViewedDate	sf_last_viewed_date	Custom	String

Contact

Amazon Q supports crawling [Salesforce Online Contacts](#) and offers the following contact field mappings.

Salesforce field name	Index field name	Description	Data type
category	_category	Default	String
sourceUrl	_source_uri	Default	String
authors	_authors	Default	String list
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date
lastModifiedBy	sf_last_modified_by	Custom	String
lastActivityDate	sf_last_activity_date	Custom	Date
createdBy	sf_created_by	Custom	String
contactName	sf_contact_name	Custom	String
phone	sf_phone	Custom	String
email	sf_email	Custom	String
department	sf_department	Custom	String
lastname	sf_lastname	Custom	String
title	sf_title	Custom	String
reportsTo	sf_reports_to	Custom	String
account	sf_account	Custom	String
otherStreet	sf_other_street	Custom	String

Salesforce field name	Index field name	Description	Data type
otherCity	sf_other_city	Custom	String
otherState	sf_other_state	Custom	String
otherPostalCode	sf_other_postal_code	Custom	String
otherCountry	sf_other_country	Custom	String
otherLatitude	sf_other_latitude	Custom	String
otherLongitude	sf_other_longitude	Custom	String
otherGeocodeAccuracy	sf_other_geocode_accuracy	Custom	String
mailingStreet	sf_mailing_street	Custom	String
mailingCity	sf_mailing_city	Custom	String
mailingState	sf_mailing_state	Custom	String
mailingPostalCode	sf_mailing_postal_code	Custom	String
mailingCountry	sf_mailing_country	Custom	String
mailingLatitude	sf_mailing_latitude	Custom	String
mailingLongitude	sf_mailing_longitude	Custom	String
mailingGeocodeAccuracy	sf_mailing_geocode_accuracy	Custom	String
fax	sf_fax	Custom	String
mobilePhone	sf_mobile_phone	Custom	String
homePhone	sf_home_phone	Custom	String

Salesforce field name	Index field name	Description	Data type
otherPhone	sf_other_phone	Custom	String
assistantPhone	sf_assistant_phone	Custom	String
assistantName	sf_assistant_name	Custom	String
leadSource	sf_lead_source	Custom	String
birthDate	sf_birthdate	Custom	Date
jigsaw	sf_jigsaw	Custom	String

Contract

Amazon Q supports crawling [Salesforce Online Contracts](#) and offers the following contract field mappings.

Salesforce field name	Index field name	Description	Data type
category	_category	Default	String
sourceUrl	_source_uri	Default	String
authors	_authors	Default	String list
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date
authors	_authors	Default	String list
accountId	_sf_accoung_id	Custom	String
ownerExpirationNotice	sf_owner_expiration_notice	Custom	String

Salesforce field name	Index field name	Description	Data type
billingStreet	sf_billing_street	Custom	String
billingCity	sf_billing_city	Custom	String
billingState	sf_billing_state	Custom	String
billingPostalCode	sf_billing_postal_code	Custom	String
billingCountry	sf_billing_country	Custom	String
contractTerm	sf_contract_term	Custom	String
ownerId	sf_owner_id	Custom	String
status	sf_status	Custom	String
customerSignedTitle	sf_customer_signed_title	Custom	String
specialTerms	sf_special_terms	Custom	String
statusCode	sf_status_code	Custom	String
contractNumber	sf_contract_number	Custom	String
lastViewedDate	sf_last_viewed_date	Custom	Date
lastReferenceDate	sf_last_reference_date	Custom	Date
billingAddressCity	sf_billing_address_city	Custom	String
billingAddressCountry	sf_billing_address_country	Custom	String

Salesforce field name	Index field name	Description	Data type
billingAddressPostalCode	sf_billing_address_postal_code	Custom	String
billingAddressState	sf_billing_address_state	Custom	String
billingAddressStreet	sf_billing_address_street	Custom	String
pricebookDescription	sf_pricebook_description	Custom	String
pricebookId	sf_pricebook_id	Custom	String
billingLatitude	sf_billing_latitude	Custom	String
billingLongitude	sf_billing_longitude	Custom	String
billingGeocodeAccuracy	sf_billing_geocode_accuracy	Custom	String
companySignedId	sf_company_signed_id	Custom	String
companySignedDate	sf_company_signed_date	Custom	Date
customerSignedId	sf_customer_signed_id	Custom	String
activatedById	sf_activated_by_id	Custom	String
activatedDate	sf_activated_date	Custom	Date
lastApprovedDate	sf_last_approved_date	Custom	Date

Salesforce field name	Index field name	Description	Data type
lastActivityDate	sf_last_activity_date	Custom	Date
accountName	sf_account_name	Custom	String
startDate	sf_start_date	Custom	Date
endDate	sf_end_date	Custom	Date
createdBy	sf_created_by	Custom	String

Document

Amazon Q supports crawling [Salesforce Online Documents](#) and offers the following document field mappings.

Salesforce field name	Index field name	Description	Data type
category	_category	Default	String
sourceUrl	_source_uri	Default	String
author	_authors	Default	String list
createdAt	_created_at	Default	Date
folder	sf_folder_name	Custom	String
isInternalUseOnly	sf_is_internal_use_only	Custom	String
isPublic	sf_is_public	Custom	String
keywords	sf_keywords	Custom	String
lastModifiedBy	sf_last_modified_by	Custom	String

Salesforce field name	Index field name	Description	Data type
updatedAt	_last_updated_at	Default	Date
fileName	sf_file_name	Custom	String
fileType	_file_type	Default	String
fileSize	sf_file_size	Custom	Long (numeric)
createdBy	sf_created_by	Custom	String
isBodySearchable	sf_is_body_searchable	Custom	String

Group

Amazon Q supports crawling [Salesforce Online Groups](#) and offers the following group field mappings.

Salesforce field name	Index field name	Description	Data type
category	_category	Default	String
sourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
groupEmail	sf_group_email	Custom	String
lastModifiedBy	sf_last_modified_by	Custom	String
lastModifiedDate	_last_modified_at	Default	Date
ownerId	sf_owner_id	Custom	String
groupName	sf_group_name	Custom	String

Salesforce field name	Index field name	Description	Data type
createdBy	_authors	Default	String list
lastFeedModifiedDate	sf_last_feed_modified_date	Custom	Date
hasPrivateFieldsAccess	sf_has_private_fields_access	Custom	String
canHaveGuests	sf_can_have_guests	Custom	String
isArchived	sf_is_archived	Custom	String
isAutoArchived	sf_is_auto_archive_disabled	Custom	String
memberCount	sf_member_count	Custom	String
collaborationType	sf_collabotration_type	Custom	String
informationTitle	sf_information_title	Custom	String

Idea

Amazon Q supports crawling [Salesforce Online Ideas](#) and offers the following idea field mappings.

Salesforce field name	Index field name	Description	Data type
category	_category	Default	String
sourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
lastModifiedBy	sf_last_modified_by	Custom	String

Salesforce field name	Index field name	Description	Data type
title	sf_title	Custom	String
status	sf_status	Custom	String
createdByName	sf_created_by	Custom	String
parentIdea	sf_parent_idea_id	Custom	String
parentIdeaId	sf_parent_idea_id	Custom	String
lastModifiedDate	_last_modified_at	Default	Date
recordTypeId	sf_record_type_id	Custom	String
communityId	sf_community_id	Custom	String
numComments	sf_number_of_comments	Custom	Long (numeric)
voteScore	sf_vote_score	Custom	Long (numeric)
voteTotal	sf_vote_total	Custom	Long (numeric)
lastCommentDate	sf_last_comment_date	Custom	Date

Lead

Amazon Q supports crawling [Salesforce Online Leads](#) and offers the following lead field mappings.

Salesforce field name	Index field name	Description	Data type
category	_category	Default	String
sourceUrl	_source_uri	Default	String

Salesforce field name	Index field name	Description	Data type
city	sf_city	Custom	String
company	sf_company	Custom	String
country	sf_country	Custom	String
createdAt	_created_at	Default	Date
lastModifiedBy	sf_last_modified_by	Custom	String
updatedAt	_last_updated_at	Default	Date
leadSource	sf_lead_source	Custom	String
state	sf_state	Custom	String
status	sf_status	Custom	String
convertedAccount	sf_converted_account	Custom	String
convertedAccountId	sf_converted_account_id	Custom	String
convertedContact	sf_converted_contact	Custom	String
convertedContactId	sf_converted_contact_id	Custom	String
convertedDate	sf_converted_date	Custom	Date
convertedOpportunity	sf_converted_opportunity	Custom	String
convertedOpportunityId	sf_converted_opportunity_id	Custom	String
firstName	sf_first_name	Custom	String

Salesforce field name	Index field name	Description	Data type
createdBy	_authors	Default	String list
isConverted	sf_is_converted	Custom	String
owner	sf_owner_name	Custom	String
lastActivityDate	sf_last_activity_date	Custom	Date
ownerId	sf_owner_id	Custom	String
lastName	sf_last_name	Custom	String
title	sf_title	Custom	String
street	sf_street	Custom	String
postalCode	sf_postal_code	Custom	String
latitude	sf_latitude	Custom	String
longitude	sf_longitude	Custom	String
geocodeAccuracy	sf_geocode_accuracy	Custom	String
phone	sf_phone	Custom	String
email	sf_email	Custom	String
industry	sf_industry	Custom	String
rating	sf_rating	Custom	String
annualRevenue	sf_annual_revenue	Custom	String
numberOfEmployees	sf_number_of_employees	Custom	Long (numeric)
jigsaw	sf_jigsaw	Custom	String

Salesforce field name	Index field name	Description	Data type
jigsawContactId	sf_jigsaw_contact_id	Custom	String
emailBouncedReason	sf_email_bounced_reason	Custom	String
individualId	sf_individual_id	Custom	String

Opportunity

Amazon Q supports crawling [Salesforce Online Opportunities](#) and offers the following opportunity field mappings.

Salesforce field name	Index field name	Description	Data type
category	_category	Default	String
sourceUrl	_source_uri	Default	String
accountName	sf_account_name	Custom	String
amount	sf_amount	Custom	String
campaign	sf_campaign_name	Custom	String
createdAt	_created_at	Default	Date
createdBy	sf_created_by	Custom	String
lastModifiedBy	sf_last_modified_by	Custom	String
lastModifiedDate	_last_updated_at	Default	Date
fiscalQuarter	sf_fiscal_quarter	Custom	String
fiscalYear	sf_fiscal_year	Custom	String

Salesforce field name	Index field name	Description	Data type
isClosed	sf_is_closed	Custom	String
isWon	sf_is_won	Custom	String
leadSource	sf_lead_source	Custom	String
opportunityName	sf_opportunity_name	Custom	String
accountId	sf_account_id	Custom	String
campaignId	sf_campaign_id	Custom	String
closeDate	sf_close_date	Custom	Date
typeValue	sf_type_value	Custom	String
lastActivityDate	sf_last_activity_date	Date	String
ownerName	sf_owner_name	Custom	String
ownerId	sf_owner_id	Custom	String
stageName	sf_stage_name	Custom	String
probablity	sf_probability	Custom	Long (numeric)
nextStep	sf_next_step	Custom	String
forestCategory	sf_forecast_category	Custom	String
forestCategoryName	sf_forecast_category_name	Custom	String
hasOpportunityLine Item	sf_has_opportunity_line_item	Custom	String
pricebook2id	sf_pricebook2_id	Custom	String
pushCount	sf_push_count	Custom	String

Salesforce field name	Index field name	Description	Data type
fiscal	sf_fiscal	Custom	String
contactId	sf_contact_id	Custom	String
lastViewedDate	sf_last_viewed_date	Custom	Date
hasOpenActivity	sf_has_open_activity	Custom	Long (numeric)
hasOverdueTask	sf_has_overdue_task	Custom	String

Partner

Amazon Q supports crawling [Salesforce Online Partner](#) and offers the following partner field mappings.

Salesforce field name	Index field name	Description	Data type
category	_category	Default	String
sourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date
createdBy	_authors	Default	String list
opportunityId	sf_opportunity_id	Custom	String
accountFromId	sf_account_from_id	Custom	String
accountTold	sf_role	Custom	String
role	sf_role	Custom	String
isPrimary	sf_is_primary	Custom	String

Salesforce field name	Index field name	Description	Data type
systemModstamp	sf_system_modstamp	Custom	Date
reversePartnerId	sf_reverse_partner_id	Custom	String
opportunity	sf_opportunity	Custom	String
accountFrom	sf_account_from	Custom	String
accountTo	sf_account_to	Custom	String

Pricebook

Amazon Q supports crawling [Salesforce Online Pricebooks](#) and offers the following pricebook field mappings.

Salesforce field name	Index field name	Description	Data type
category	_category	Default	String
sourceUrl	_source_uri	Default	String
isActive	sf_is_active	Custom	String
lastModifiedBy	sf_last_modified_by	Default	String
lastModifiedDate	_last_updated_at	Default	Date
pricebookName	sf_pricebook_name	Custom	String
createdAt	_created_at	Default	Date
createdBy	_authors	Default	String list
lastViewedDate	sf_last_viewed_date	Custom	Date
isStandard	sf_is_standard	Custom	String

Product

Amazon Q supports crawling [Salesforce Online Product](#) and offers the following product field mappings.

Salesforce field name	Index field name	Description	Data type
category	_category	Default	String
sourceUrl	_source_uri	Default	String
family	sf_family	Custom	String
isActive	sf_is_active	Custom	String
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date
lastModifiedBy	sf_last_modified_by	Custom	String
productCode	sf_product_code	Custom	String
createdBy	_authors	Default	String list
productName	sf_product_name	Custom	String
externalDataSourceId	sf_external_data_source_id	Custom	String
externalId	sf_external_id	Custom	String
displayUrl	sf_display_url	Custom	String
quantityUnitOfMeasure	sf_quantity_unit_of_measure	Custom	String
isArchived	sf_is_archived	Custom	String
lastViewedDate	sf_last_viewed_date	Custom	Date

Salesforce field name	Index field name	Description	Data type
stockKeepingUnit	sf_stock_keeping_unit	Custom	String

Solution

Amazon Q supports crawling [Salesforce Online Solutions](#) and offers the following solution field mappings.

Salesforce field name	Index field name	Description	Data type
category	_category	Default	String
sourceUrl	_source_uri	Default	String
isPublished	sf_is_published	Custom	String
isReviewed	sf_is_reviewed	Custom	String
lastModifiedBy	sf_last_modified_by	Custom	String
lastModifiedDate	_last_updated_at	Default	Date
ownerName	sf_owner_name	Custom	String
solutionNumber	sf_solution_number	Custom	String
status	sf_status	Custom	String
timesUsed	sf_times_used	Custom	String
solutionName	sf_solution_name	Custom	String
createdByName	_authors	Default	String list
createdAt	_created_at	Default	Date

Salesforce field name	Index field name	Description	Data type
solutionNote	sf_solution_note	Custom	String
ownderId	sf_ownderId	Custom	String
lastViewedDate	sf_last_viewed_date	Custom	Date

Profile

Amazon Q supports crawling [Salesforce Online Profiles](#) and offers the following profile field mappings.

Salesforce field name	Index field name	Description	Data type
category	_category	Default	String
sourceUrl	_source_uri	Default	String
updatedAt	_last_updated_at	Default	Date
lastModifiedBy	sf_last_modified_by	Custom	String
createdBy	_authors	Default	String list
createdAt	_created_at	Default	Date
userType	sf_user_type	Custom	String

Task

Amazon Q supports crawling [Salesforce Online Tasks](#) and offers the following task field mappings.

Salesforce field name	Index field name	Description	Data type
category	_category	Default	String
sourceUrl	_source_uri	Default	String
accountName	sf_account_name	Custom	String
lastModifiedBy	sf_last_modified_by	Custom	String
lastModifiedDate	_last_updated_at	Default	Date
ownerName	sf_owner_name	Custom	String
isRecurrence	sf_is_recurrence	Custom	String
isClosed	sf_is_closed	Custom	String
isArchived	sf_is_archived	Custom	String
priority	sf_priority	Custom	String
status	sf_status	Custom	String
whatId	sf_what_id	Custom	String
createdByName	_authors	Default	String list
createdAt	_created_at	Default	Date
subject	sf_subject	Custom	String
activityDate	sf_activity_date	Custom	Date
activityDate	sf_activity_date	Custom	Date
isHighPriority	sf_is_high_priority	Custom	String
ownerId	sf_owner_id	Custom	String
callType	sf_call_type	Custom	String

User

Amazon Q supports crawling [Salesforce Online Users](#) and offers the following user field mappings.

Salesforce field name	Index field name	Description	Data type
category	_category	Default	String
sourceUrl	_source_uri	Default	String
account	sf_account	Custom	String
isActive	sf_is_active	Custom	String
city	sf_city	Custom	String
lastModifiedBy	sf_last_modified_by	Custom	String
updatedAt	_last_updated_at	Default	Date
companyName	sf_company_name	Custom	String
country	sf_country	Custom	String
department	sf_department	Custom	String
division	sf_division	Custom	String
email	sf_email	Custom	String
employeeNumber	sf_employee_number	Custom	String
firstName	sf_first_name	Custom	String
lastName	sf_last_name	Custom	String
manager	sf_manager	Custom	String
state	sf_state	Custom	String
userRole	sf_user_role	Custom	String

Salesforce field name	Index field name	Description	Data type
username	sf_username	Custom	String
createdBy	_authors	Default	String list
createdAt	_created_at	Default	Date
street	sf_street	Custom	String
postalCode	sf_postal_code	Custom	String
latitude	sf_latitude	Custom	String
longitude	sf_longitude	Custom	String
geocodeAccuracy	sf_geocode_accuracy	Custom	String
phone	sf_phone	Custom	String
fax	sf_fax	Custom	String
mobilePhone	sf_mobile_phone	Custom	String
profileName	sf_profile_name	Custom	String
aboutMe	sf_about_me	Custom	String
languageLocaleKey	sf_language_locale_key	Custom	String

Chatter

Amazon Q supports crawling [Salesforce Online Chatters](#) and offers the following chatter field mappings.

Salesforce field name	Index field name	Description	Data type
category	_category	Default	String
sourceUrl	_source_uri	Default	String
body	sf_body	Custom	String
createdAt	_created_at	Default	Date
lastEditById	sf_last_edit_by_id	Custom	String
lastEditDate	sf_last_edit_date	Custom	Date
lastModifiedDate	_last_updated_at	Default	Date
insertedById	sf_inserted_by_id	Custom	String
createdBy	_authors	Default	String list
parentId	sf_parent_id	Custom	String
revision	sf_revision	Custom	String
status	sf_status	Custom	String
isRichText	sf_is_rich_text	Custom	String

Knowledge articles

Amazon Q supports crawling [Salesforce Online Knowledge articles](#) and offers the following knowledge article field mappings.

Salesforce field name	Index field name	Description	Data type
category	_category	Default	String
sourceUrl	_source_uri	Default	String

Salesforce field name	Index field name	Description	Data type
articleTitle	sf_title	Custom	String
articleNumber	sf_article_number	Default	Date
knowledgeArticleId	sf_knowledge_article_id	Custom	String
lastPublishedDate	sf_last_published_date	Custom	Date
publishStatus	sf_publish_status	Custom	String
versionNumber	sf_version_number	Custom	String
language	sf_language	Custom	String
ownerId	sf_owner_id	Custom	String
summary	sf_summary	Custom	String
firstPublishedDate	sf_first_published_date	Custom	Date
updatedAt	_last_updated_at	Default	Date
archivedDate	sf_archived_date	Custom	Date
isLatestVersion	sf_is_latest_version	Custom	String
sourceId	sf_sourceId	Custom	String
createdBy	_authors	Default	String list
assignmentDate	sf_assignment_date	Custom	Long (numeric)
assignmentDueDate	sf_assignment_due_date	Custom	Date

Salesforce field name	Index field name	Description	Data type
articleCaseAttachCount	sf_article_case_attach_count	Custom	Long (numeric)
articleTotalViewCount	sf_article_total_view_count	Custom	Long (numeric)
urlName	sf_url_name	Custom	String
assignmentNote	sf_assignment_date	Custom	String
migratedToFromArticleVersion	sf_migrated_article_version	Custom	String
assignedBy	sf_assigned_by	Custom	String
assignedTo	sf_assigned_to	Custom	Date

Attachments

Amazon Q supports crawling [Salesforce Online Attachments](#) and offers the following attachment field mappings.

Salesforce field name	Index field name	Description	Data type
category	_category	Default	String
sourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date
fileName	sf_file_name	Custom	String
fileType	_file_type	Default	String

Salesforce field name	Index field name	Description	Data type
fileSize	sf_file_size	Custom	Long (numeric)
parentName	sf_parent_name	Default	String
createdBy	_authors	Default	String list

Custom object

Amazon Q supports crawling custom objects and offers the following custom object field mappings.

Salesforce field name	Index field name	Description	Data type
category	_category	Default	String
sourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date
lastModifiedById	sf_last_modified_by_id	Custom	String
customObjectName	sf_custom_object_name	Custom	String
createdBy	_authors	Default	String list
lastModifiedBy	sf_last_modified_by	Custom	String
documentbody	_document_body	Custom	String

IAM role for Amazon Q Business Salesforce Online connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the `BatchPutDocument` and `BatchDeleteDocument` operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToGetSecret",
      "Effect": "Allow",
      "Action": [
        "secretsmanager:GetSecretValue"
      ],
      "Resource": [
        "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
      ]
    },
    {
      "Sid": "AllowsAmazonQToDecryptSecret",
      "Effect": "Allow",
      "Action": [
        "kms:Decrypt"
      ],
      "Resource": [
```

```

    "arn:aws:kms:{{region}}:{{account_id}}:key/[{{key_id}}]"
  ],
  "Condition": {
    "StringLike": {
      "kms:ViaService": [
        "secretsmanager.*.amazonaws.com"
      ]
    }
  }
},
{
  "Sid": "AllowsAmazonQToIngestDocuments",
  "Effect": "Allow",
  "Action": [
    "qbusiness:BatchPutDocument",
    "qbusiness:BatchDeleteDocument"
  ],
  "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
},
{
  "Sid": "AllowsAmazonQToIngestPrincipalMapping",
  "Effect": "Allow",
  "Action": [
    "qbusiness:PutGroup",
    "qbusiness:CreateUser",
    "qbusiness>DeleteGroup",
    "qbusiness:UpdateUser",
    "qbusiness:ListGroups"
  ],
  "Resource": [
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}",
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}",
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}/data-source/*"
  ]
},
{
  "Sid": "AllowsAmazonQToCreateAndDeleteNI",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateNetworkInterface",
    "ec2>DeleteNetworkInterface"
  ]
}

```

```

    ],
    "Resource": [
      "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
      "arn:aws:ec2:{{region}}:{{account_id}}:security-group/[{{security_group}}]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2:DeleteNetworkInterface"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:RequestTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
      },
      "ForAllValues:StringEquals": {
        "aws:TagKeys": [
          "AMAZON_Q"
        ]
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateTags",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateTags"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringEquals": {
        "ec2:CreateAction": "CreateNetworkInterface"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterfacePermission"
    ],
  },

```



```

    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:ResourceTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
    "Effect": "Allow",
    "Action": [
      "ec2:DescribeNetworkInterfaces",
      "ec2:DescribeAvailabilityZones",
      "ec2:DescribeNetworkInterfaceAttribute",
      "ec2:DescribeVpcs",
      "ec2:DescribeRegions",
      "ec2:DescribeNetworkInterfacePermissions",
      "ec2:DescribeSubnets"
    ],
    "Resource": "*"
  }
]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "{{source_account}}"
        },
        "ArnEquals": {
          "aws:SourceArn": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/{{application_id}}"
        }
      }
    }
  ]
}

```

```

    }
  }
}
]
}

```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Known limitations for the Amazon Q Business Salesforce Online connector

The Amazon Q Business Salesforce Online connector has the following known limitations:

- Salesforce Online API doesn't provide the status of deleted **Group**, **Partner**, **Profile**, and **User** entities. So, the Salesforce Online connector can't retrieve this information.
- Salesforce Online API doesn't provide the status of modified **Attachment titles** (Lightning Version). So, the Salesforce Online connector can't retrieve this information.
- Salesforce Online connector supports custom field mappings only for the following entities: **Account**, **Campaign**, **Contact**, **Contract**, **Case**, **Product Lead**, **Pricebook**, and **CustomEntity**.
- Salesforce Online API does not provide ACL information for documents with shared access types.
- By default, Salesforce Online Developer has a maximum limit of 15000 total calls per 24 hour period. If a request exceeds this limit, the API returns a REQUEST_LIMIT_EXCEEDED error.

Troubleshooting your Amazon Q Business Salesforce Online connector

The following table provides information about error codes you may see for the Salesforce Online connector and suggested troubleshooting actions.

Error code	Error message	Suggested resolution
SF-5001	Invalid HostURL.	Provide valid HostURL.
SF-5002	Invalid userName or password.	Provide valid userName or password.
SF-5003	Invalid clientSecret.	Provide valid clientSecret.
SF-5004	Invalid clientId.	Provide valid clientId.

Error code	Error message	Suggested resolution
SF-5005	Invalid grant type.	Provide valid grant type.
SF-5006	Error while generating Access Token.	Provide valid credentials or try again later.
SF-5100	Null/empty HostUrl.	Provide HostUrl.
SF-5101	Null/empty client ID.	Provide client ID.
SF-5102	Null/empty client secret	Provide client secret.
SF-5103	Null/empty user name.	Provide user name.
SF-5104	Null/empty password.	Provide password.
SF-5150	Null/empty authentication URL.	Provide authentication URL.
SF-5151	Invalid HostURL pattern.	Provide valid HostURL pattern.
SF-5152	Invalid Authentication URL.	Provide valid Authentication URL.
SF-5500	ContinuableInternalServerError.	Try again later.

Connecting ServiceNow Online to Amazon Q Business

ServiceNow provides a cloud-based service management system to create and manage organization-level workflows, such as IT services, ticketing systems, and support. You can connect ServiceNow Online instance to Amazon Q Business—using either the AWS Management Console or the [CreateDataSource](#) API—and create an Amazon Q web experience.

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).

- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [ServiceNow Online connector overview](#)
- [Prerequisites for connecting Amazon Q Business to ServiceNow Online](#)
- [Connecting Amazon Q Business to ServiceNow Online using the console](#)
- [Connecting Amazon Q Business to ServiceNow using APIs](#)
- [How Amazon Q Business connector crawls ServiceNow ACLs](#)
- [Amazon Q Business ServiceNow Online data source connector field mappings](#)
- [IAM role for Amazon Q Business ServiceNow Online connector](#)
- [Known limitations for the Amazon Q Business ServiceNow Online connector](#)
- [Troubleshooting your Amazon Q Business ServiceNow Online connector](#)

ServiceNow Online connector overview

The following table gives an overview of the Amazon Q Business ServiceNow Online connector and its supported features.

Category	Feature	Support
Security	Authentication type	Basic, OAuth 2.0
	Authentication credentials	<p>Basic</p> <ul style="list-style-type: none"> • ServiceNow Online host URL • User name • Password <p>OAuth 2.0</p> <ul style="list-style-type: none"> • ServiceNow Online host URL • User name • Password

Category	Feature	Support
		<ul style="list-style-type: none"> Client ID Client secret <div style="border: 1px solid #f08080; border-radius: 10px; padding: 10px; margin-top: 10px;">  Important Admin privileges required. </div>
	Supported versions	San Diego, Tokyo, Rome, Vancouver, Others
	<u>Access Control List (ACL) crawling</u>	Yes. For more information, see <u>ACL crawling</u> .
	<u>Identity crawling</u>	Yes
	<u>VPC</u>	Yes
Crawl features	Custom metadata	Yes. Supports custom fields for all entities.
	Entities	Yes. The following entities are supported: <ul style="list-style-type: none"> Knowledge article Knowledge article attachments Service catalog Service catalog attachments Incident Incident attachments
	<u>Field mappings</u>	Yes. Supports both default and custom field mappings. For more information, see <u>Field mappings</u> .

Category	Feature	Support
	Filters	<p>Yes. The following filters are supported :</p> <ul style="list-style-type: none"> • Crawl public knowledge articles • Crawl knowledge articles with filter query • Crawl knowledge article attachments • Use regex filters for knowledge articles • Crawl service catalog items • Crawl service catalog item attachments • Use regex filters for service catalog items • Crawl incidents • Crawl incident attachments • Crawl incidents with filter query • Use regex filters for active and inactive incidents • Including and excluding content by file type • Including and excluding content based on file name • Crawl ACL for knowledge article, service catalogs, and incidents
	<u>Sync mode</u>	Supports full and incremental sync.
	<u>File types</u>	Supports all files supported by Amazon Q.

Prerequisites for connecting Amazon Q Business to ServiceNow Online

Before you begin, make sure that you have completed the following prerequisites.

In ServiceNow, make sure you have:

- Created a Personal or Enterprise Developer Instance and have a ServiceNow instance with an administrative role.
- Copied the host of your ServiceNow instance URL. The format for the host URL you enter is *your-domain.service-now.com*. You need your ServiceNow instance URL to connect to Amazon Q.
- Configured basic authentication credentials containing a username and password to allow Amazon Q to connect to your ServiceNow instance.
- **Optional:** Configured an OAuth 2.0 credential token that can identify Amazon Q using a username, password, and a generated client ID, and a client secret. For more information, see [ServiceNow documentation on OAuth 2.0 authentication](#) on the ServiceNow website.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your ServiceNow Online authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Connecting Amazon Q Business to ServiceNow Online using the console

The following procedure outlines how to connect Amazon Q Business to ServiceNow Online using the AWS Management Console.

Connecting Amazon Q to ServiceNow Online

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **ServiceNow Online** page, enter the following information:
6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. In **Source**, enter the following information:
 - **ServiceNow host** – Enter your ServiceNow host name without the protocol. For example, *example.service-now.com*.
 - **ServiceNow version** – Select your ServiceNow version, whether **Tokyo**, **San Diego**, **Rome**, **Vancouver**, and **Others**.
8. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

Note

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

9. **Authentication** – Choose between **Basic authentication** and **OAuth 2.0 authentication** and then enter the following information for your **AWS Secrets Manager secret**.
 - a. **Secret name** – A name for your secret.
 - b. **Basic Authentication** – Enter the **Secret name**, **Username**, and **Password** for your ServiceNow account.

If using OAuth2 Authentication – Enter the **Secret name, Username, Password, Client ID,** and **Client Secret** that you created in your ServiceNow account.

c. Choose **Save and add secret**.

10. **Configure VPC and security group – optional** – Choose whether you want to use a VPC. If you do, enter the following information:

- a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
- b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

11. **Identity crawler** – Choose to activate Amazon Q identity crawler to sync identity information. For more information, see [Identity crawler](#).

12. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

13. **Sync scope** – Set the content that you want to sync.

a. For **Knowledge articles**, choose from the following options :

- **Knowledge articles** – Choose to index knowledge articles.
- **Knowledge article attachments** – Choose to index knowledge article attachments.
- **Type of knowledge articles** – Choose between **Only public articles** and **Knowledge articles based on ServiceNow filter query**, based on your use case. If you select **Include articles based on ServiceNow filter query**, you must enter a **Filter query** copied from your ServiceNow account. Example filter queries include: *workflow_state=draft^EQ, kb_knowledge_base=dfc19531bf2021003f07e2c1ac0739ab^text ISNOTEMPTY^EQ*, and *article_type=text^active=true^EQ*.

⚠ Important

If you choose to crawl **Only public articles**, Amazon Q crawls only knowledge articles assigned a public access role in ServiceNow Online.

- **Include articles based on short description filter** – Specify regular expression patterns to include or exclude specific articles.
- b. For **Service catalog items**:
- **Service catalog items** – Choose to index service catalog items.
 - **Service catalog item attachments** – Choose to index service catalog item attachments.
 - **Active service catalog items** – Choose to index active service catalog items.
 - **Inactive service catalog items** – Choose to index inactive service catalog items.
 - **Filter query** – Choose to include service catalog items based on a filter defined in your ServiceNow instance. Example filter queries include:
short_descriptionLIKEAccess^category=2809952237b1300054b6a3549dbe5dd4^EQ
nameSTARTSWITHService^active=true^EQ.
 - **Include service catalog items based on short description filter** – Specify a regex pattern to include specific catalog items.
- c. For **Incidents**:
- **Incidents** – Choose to index service incidents.
 - **Incident attachments** – Choose to index incident attachments.
 - **Active incidents** – Choose to index active incidents.
 - **Inactive incidents** – Choose to index inactive incidents.
 - **Active incident type** – Choose between **All incidents**, **Open incidents**, **Open - unassigned incidents**, and **Resolved incidents**, depending on your use case.
 - **Filter query** – Choose to include incidents based on a filter defined in your ServiceNow instance. Example filter queries include:
short_descriptionLIKETest^urgency=3^state=1^EQ, and
priority=2^category=software^EQ .
 - **Include incidents based on short description filter** – Specify a regex pattern to include specific incidents.
- d. In **Additional configuration – optional**:

- **ACL information** – Access control lists for entities that you have selected are included by default. Deselecting an access control list will make all files in that category public. ACL options are automatically deactivated for entities that aren't selected. For public articles, ACL isn't applied.
 - For **Maximum file size** – Specify the file size limit in MBs that Amazon Q will crawl. Amazon Q will crawl only the files within the size limit you define. The default file size is 50MB. The maximum file size should be greater than 0MB and less than or equal to 50MB.
 - **Attachment regex patterns** – Add regular expression patterns to include or exclude specific attached files of catalogs, knowledge articles, and incidents. You can add up to 100 patterns.
14. For **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.
- **Full sync** – Sync all content regardless of the previous sync status.
 - **New, modified, or deleted content sync** – Only sync new, modified, and deleted content.
15. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
16. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
17. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:
- a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
 - b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

 **Note**

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

18. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

19. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

Note

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to ServiceNow using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the configuration parameter to provide a JSON schema with all other configuration information specific to your data source connector.

For an example of the API request, see [CreateDataSource](#) in the Amazon Q API Reference.

ServiceNow JSON schema

The following is the ServiceNow JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
      "type": "object",
      "properties": {
        "repositoryEndpointMetadata": {
          "type": "object",
```

```

    "properties": {
      "hostUrl": {
        "type": "string",
        "pattern": "^(?!^(https?|ftp|file):\\|\\/))[a-z0-9-]+(.service-
now.com|.servicenowservices.com)$",
        "minLength": 1,
        "maxLength": 2048
      },
      "authType": {
        "type": "string",
        "enum": [
          "basicAuth",
          "OAuth2"
        ]
      },
      "servicenowInstanceVersion": {
        "type": "string",
        "enum": [
          "Tokyo",
          "SanDiego",
          "Rome",
          "Vancouver",
          "Others"
        ]
      }
    },
    "required": [
      "hostUrl",
      "authType",
      "servicenowInstanceVersion"
    ]
  },
  "required": [
    "repositoryEndpointMetadata"
  ],
  "repositoryConfigurations": {
    "type": "object",
    "properties": {
      "knowledgeArticle": {
        "type": "object",
        "properties": {
          "fieldMappings": {

```

```

    "type": "array",
    "items": [
      {
        "type": "object",
        "properties": {
          "indexFieldName": {
            "type": "string"
          },
          "indexFieldType": {
            "type": "string",
            "enum": [
              "STRING",
              "DATE",
              "STRING_LIST"
            ]
          },
          "dataSourceFieldName": {
            "type": "string"
          },
          "dateFieldFormat": {
            "type": "string",
            "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
          }
        },
        "required": [
          "indexFieldName",
          "indexFieldType",
          "dataSourceFieldName"
        ]
      }
    ]
  },
  "required": [
    "fieldMappings"
  ],
  "attachment": {
    "type": "object",
    "properties": {
      "fieldMappings": {
        "type": "array",
        "items": [
          {

```

```

        "type": "object",
        "properties": {
          "indexFieldName": {
            "type": "string"
          },
          "indexFieldType": {
            "type": "string",
            "enum": [
              "STRING",
              "LONG",
              "DATE",
              "STRING_LIST"
            ]
          },
          "dataSourceFieldName": {
            "type": "string"
          },
          "dateFieldFormat": {
            "type": "string",
            "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
          }
        },
        "required": [
          "indexFieldName",
          "indexFieldType",
          "dataSourceFieldName"
        ]
      }
    ],
    "required": [
      "fieldMappings"
    ],
    "serviceCatalog": {
      "type": "object",
      "properties": {
        "fieldMappings": {
          "type": "array",
          "items": [
            {
              "type": "object",
              "properties": {

```

```

        "indexFieldName": {
          "type": "string"
        },
        "indexFieldType": {
          "type": "string",
          "enum": [
            "STRING",
            "DATE",
            "STRING_LIST"
          ]
        },
        "dataSourceFieldName": {
          "type": "string"
        },
        "dateFieldFormat": {
          "type": "string",
          "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
        }
      ],
      "required": [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
      ]
    }
  ]
}
},
"required": [
  "fieldMappings"
]
},
"incident": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": [
        {
          "type": "object",
          "properties": {
            "indexFieldName": {
              "type": "string"
            }
          }
        }
      ]
    }
  }
}

```



```

        "indexFieldType": {
          "type": "string",
          "enum": [
            "STRING",
            "DATE",
            "STRING_LIST"
          ]
        },
        "dataSourceFieldName": {
          "type": "string"
        },
        "dateFieldFormat": {
          "type": "string",
          "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
        }
      ],
      "required": [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
      ]
    }
  ]
}
},
"required": [
  "fieldMappings"
]
}
},
"additionalProperties": {
  "type": "object",
  "properties": {
    "maxFileSizeInMegaBytes": {
      "type": "string"
    },
    "isCrawlKnowledgeArticle": {
      "type": "boolean"
    },
    "isCrawlKnowledgeArticleAttachment": {
      "type": "boolean"
    },
    "includePublicArticlesOnly": {

```

```
    "type": "boolean"
  },
  "knowledgeArticleFilter": {
    "type": "string"
  },
  "incidentQueryFilter": {
    "type": "string"
  },
  "serviceCatalogQueryFilter": {
    "type": "string"
  },
  "isCrawlServiceCatalog": {
    "type": "boolean"
  },
  "isCrawlServiceCatalogAttachment": {
    "type": "boolean"
  },
  "isCrawlActiveServiceCatalog": {
    "type": "boolean"
  },
  "isCrawlInactiveServiceCatalog": {
    "type": "boolean"
  },
  "isCrawlIncident": {
    "type": "boolean"
  },
  "isCrawlIncidentAttachment": {
    "type": "boolean"
  },
  "isCrawlActiveIncident": {
    "type": "boolean"
  },
  "isCrawlInactiveIncident": {
    "type": "boolean"
  },
  "applyACLForKnowledgeArticle": {
    "type": "boolean"
  },
  "applyACLForServiceCatalog": {
    "type": "boolean"
  },
  "applyACLForIncident": {
    "type": "boolean"
  },
}
```

```
"incidentStateType": {
  "type": "array",
  "items": {
    "type": "string",
    "enum": [
      "Open",
      "Open - Unassigned",
      "Resolved",
      "All"
    ]
  }
},
"knowledgeArticleTitleRegExp": {
  "type": "string"
},
"serviceCatalogTitleRegExp": {
  "type": "string"
},
"incidentTitleRegExp": {
  "type": "string"
},
"inclusionFileTypePatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"exclusionFileTypePatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"inclusionFileNamePatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
},
"exclusionFileNamePatterns": {
  "type": "array",
  "items": {
    "type": "string"
  }
}
```

```
    }
  },
  "required": []
},
"type": {
  "type": "string",
"enum": [
  "SERVICENOWV2",
  "SERVICENOW"
]
},
"enableIdentityCrawler": {
  "type": "boolean"
},
"syncMode": {
  "type": "string",
  "enum": [
    "FORCED_FULL_CRAWL",
    "FULL_CRAWL"
  ]
},
"secretArn": {
  "type": "string",
  "minLength": 20,
  "maxLength": 2048
}
},
"version": {
  "type": "string",
  "anyOf": [
    {
      "pattern": "1.0.0"
    }
  ]
},
"required": [
  "connectionConfiguration",
  "repositoryConfigurations",
  "syncMode",
  "additionalProperties",
  "secretArn",
  "type"
]
```

}

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.
repositoryEndpointMetadata	The endpoint information for the data source.
hostUrl	The ServiceNow host URL. For example, <i>your-domain.service-now.com</i> .
authType	The type of authentication you are using, either <code>basicAuth</code> or <code>OAuth2</code> .
servicenowInstanceVersion	The ServiceNow version you are using. You can choose between Tokyo, San Diego, Rome, Vancouver, and Others.
repositoryConfigurations	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings.
<ul style="list-style-type: none"> knowledgeArticle attachment serviceCatalog incident 	A list of ServiceNow objects that Amazon Q crawls and maps the attributes of to Amazon Q index field names.
additionalProperties	Additional configuration options for your content in your data source.
maxFileSizeInMegabytes	Specify the file size limit in MBs that Amazon Q will crawl. Amazon Q will crawl only the files within the size limit you define. The default file size is 50MB. The maximum file size should be greater than 0MB and less than or equal to 50MB.

Configuration	Description
<ul style="list-style-type: none"> • knowledgeArticleFilter • incidentQueryFilter • serviceCatalogQueryFilter 	Specify specific knowledge articles, incident queries, and service catalog queries to crawl.
incidentStateType	Specify incidents to crawl by state type: whether Open, Open - Unassigned Resolved or All.
<ul style="list-style-type: none"> • knowledgeArticleTitleRegExp • serviceCatalogTitleRegExp • incidentTitleRegExp • inclusionFileTypePatterns • exclusionFileTypePatterns • inclusionFileNamePatterns • exclusionFileNamePatterns 	A list of regular expression patterns to include and exclude specific files in your ServiceNow data source. Files that match the patterns are included in the index. Files that don't match the patterns are excluded from the index. If a file matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence and the file isn't included in the index.

Configuration	Description
<ul style="list-style-type: none"> • <code>isCrawlKnowledgeArticle</code> • <code>isCrawlKnowledgeArticleAttachment</code> • <code>includePublicArticlesOnly</code> • <code>isCrawlServiceCatalog</code> • <code>isCrawlServiceCatalogAttachment</code> • <code>isCrawlActiveServiceCatalog</code> • <code>isCrawlInactiveServiceCatalog</code> • <code>isCrawlIncident</code> • <code>isCrawlIncidentAttachment</code> • <code>isCrawlActiveIncident</code> • <code>isCrawlInactiveIncident</code> • <code>applyACLForKnowledgeArticle</code> • <code>applyACLForServiceCatalog</code> • <code>applyACLForIncident</code> 	<p><code>true</code> to index ServiceNow knowledge articles, service catalogs, incidents, and attachments and their ACLs.</p>
<p><code>type</code></p>	<p>The type of data source. Specify <code>SERVICENOWV2</code> as your data source type.</p>
<p><code>enableIdentityCrawler</code></p>	<p><code>true</code> to activate identity crawler. Identity crawler is activated by default. Crawling identity information on users and groups with access to specific documents is useful for user context filtering. Search results are filtered based on the user or their group access to documents. See Identity crawler for more information.</p>

Configuration	Description
syncMode	<p>Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose between the following options:</p> <ul style="list-style-type: none">• Use <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index• Use <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index
secretARN	<p>The Amazon Resource Name (ARN) of an AWS Secrets Manager secret that contains the key-value pairs required to connect to your ServiceNow . The secret must contain a JSON structure with the following keys:</p> <pre data-bbox="699 993 1507 1192">{ "username": "user name", "password": "password" }</pre> <p>If you use OAuth2 authentication, your secret must contain a JSON structure with the following keys:</p> <pre data-bbox="699 1346 1507 1623">{ "username": "user name", "password": "password", "clientId": "client id", "clientSecret": "client secret" }</pre>
version	The version of the template that's currently supported.

How Amazon Q Business connector crawls ServiceNow ACLs

When you connect an ServiceNow data source to Amazon Q Business, Amazon Q Business crawls ACL information attached to a document (user and group information) from your ServiceNow instance. If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

The group and user IDs are mapped as follows:

- `_group_ids` – Group IDs exist in ServiceNow on files where there are set access permissions. They're mapped from the role names of `sys_ids` in ServiceNow .
- `_user_id` – User IDs exist in ServiceNow on files where there are set access permissions. They're mapped from the user emails as the IDs in ServiceNow .

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business ServiceNow Online data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q Business enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a attribute mapping already available, or if you want to map additional document attributes to index fields, use the custom field mappings to specify how a data source

attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q ServiceNow connector supports the following entities and the associated reserved and custom attributes.

Supported entities and field mappings

- [Knowledge articles](#)
- [Service catalog](#)
- [Attachments](#)
- [Incidents](#)

Knowledge articles

Amazon Q supports crawling [ServiceNow Online Knowledge articles](#) and offers the following knowledge article field mappings.

ServiceNow field name	Index field name	Description	Data type
text	sn_ka_text	Custom	String
short_description	sn_ka_short_description	Custom	String
sys_created_on	_created_at	Default	Date
sys_updated_on	_last_updated_at	Default	Date

ServiceNow field name	Index field name	Description	Data type
kb_category_name	_category	Default	String
sys_created_by	_authors	Default	String
sys_updated_by	sn_updatedBy	Custom	String
sys_id	sn_sys_id	Custom	String
published	sn_ka_publish_date	Custom	Date
workflow_state	sn_ka_workflow_state	Custom	String
kb_category	sn_ka_category	Custom	String
article_type	sn_ka_article_type	Custom	String
first_name	sn_ka_first_name	Custom	String
last_name	sn_ka_last_name	Custom	String
user_name	sn_ka_user_name	Custom	String
valid_to	sn_ka_valid_to	Custom	Date
kb_knowledge_base	sn_ka_knowledge_base	Custom	String
number	sn_ka_number	Custom	String
url	sn_url	Custom	String
displayUrl	_source_uri	Default	String
replacement_article	sn_ka_replacement_article	Custom	String
description	sn_ka_description	Custom	String

ServiceNow field name	Index field name	Description	Data type
wiki	sn_ka_wiki	Custom	String
rating	sn_ka_rating	Custom	String
rating	sn_ka_rating	Custom	String
view_as_allowed	sn_ka_view_as_allowed	Custom	String
source	sn_ka_source	Custom	String
image	sn_ka_image	Custom	String
author	sn_ka_author	Custom	String
active	sn_ka_active	Custom	String
helpful_count	sn_ka_helpful_count	Custom	String
meta_description	sn_ka_meta_description	Custom	String
meta	sn_ka_meta	Custom	String
topic	sn_ka_topic	Custom	String
roles	sn_ka_roles	Custom	String
disable_suggesting	sn_ka_disable_suggesting	Custom	String
use_count	sn_ka_use_count	Custom	String
flagged	sn_ka_flagged	Custom	String
disable_commenting	sn_ka_disable_commenting	Custom	String

ServiceNow field name	Index field name	Description	Data type
retired	sn_ka_retired	Custom	String
display_attachments	sn_ka_display_attachments	Custom	String
taxonomy_topic	sn_ka_taxonomy_topic	Custom	String

Service catalog

Amazon Q supports crawling [ServiceNow Online service catalogs](#) and offers the following service catalog field mappings.

ServiceNow field name	Index field name	Description	Data type
description	sn_sc_description	Custom	String
short_description	sn_sc_short_description	Custom	String
sys_created_on	_created_at	Default	Date
sys_updated_on	_last_updated_at	Default	Date
category_name	_category	Default	String
sys_created_by	_authors	Default	String list
sys_updated_by	sn_updated_by	Custom	String
sys_id	sn_sys_id	Custom	String
sc_catalogs	sn_sc_catalogs	Custom	String
sc_catalogs_name	sn_sc_catalogs_name	Custom	String

ServiceNow field name	Index field name	Description	Data type
category	sn_sc_category	Custom	String
category_full_name	sn_sc_category	Custom	String
url	sn_url	Custom	String
displayUrl	_source_uri	Default	String
show_variable_help_on_load	sn_sc_show_var_help_on_load	Custom	String
no_order_now	sn_sc_no_order_now	Custom	String
sc_ic_version	sn_sc_sc_ic_version	Custom	String
delivery_time	sn_sc_deliver_time	Custom	String
published_ref	sn_sc_published_ref	Custom	String
price	sn_sc_price	Custom	String
recurring_frequency	sn_sc_recurring_frequency	Custom	String
sys_name	sn_sc_sys_name	Custom	String
model	sn_sc_model	Custom	String
state	sn_sc_state	Custom	String
no_cart	sn_sc_no_cart	Custom	String
group	sn_sc_group	Custom	String
hide_sp	sn_sc_hide_sp	Custom	String
order	sn_sc_order	Custom	String
start_closed	sn_sc_start_closed	Custom	String

ServiceNow field name	Index field name	Description	Data type
image	sn_sc_image	Custom	String
no_quantity	sn_sc_no_quantity	Custom	String
delivery_plan	sn_sc_delivery_plan	Custom	String
active	sn_sc_active	Custom	String
checked_out	sn_sc_checked_out	Custom	String
custom_cart	sn_sc_custom_cart	Custom	String
no_cart_v2	sn_sc_no_cart_v2	Custom	String
no_proceed_checkout	sn_sc_no_proceed_checkout	Custom	String
ignore_price	sn_sc_ignore_price	Custom	String
sys_update_name	sn_sc_sys_update_name	Custom	String
meta	sn_sc_meta	Custom	String
omit_price	sn_sc_omit_price	Custom	String
name	sn_sc_name	Custom	String
mobile_hide_price	sn_sc_mobile_hide_price	Custom	String
no_wishlist_v2	sn_sc_no_wishlist_v2	Custom	String
preview	sn_sc_preview	Custom	String
type	sn_sc_type	Custom	String
access_type	sn_sc_access_type	Custom	String

ServiceNow field name	Index field name	Description	Data type
roles	sn_sc_roles	Custom	String
icon	sn_sc_icon	Custom	String
mobile_picture	sn_sc_mobile_picture	Custom	String
availability	sn_sc_availability	Custom	String
mandatory_attachment	sn_sc_mandatory_attachment	Custom	String
request_method	sn_sc_request_method	Custom	String
visible_guide	sn_sc_visible_guide	Custom	String
visible_standalone	sn_sc_visible_standalone	Custom	String
no_order	sn_sc_no_order	Custom	String
vendor	sn_sc_vendor	Custom	String
no_attachment_v2	sn_sc_no_attachment_v2	Custom	String
mobile_picture_type	sn_sc_mobile_picture_type	Custom	String
visible_bundle	sn_sc_visible_bundle	Custom	String
ordered_item_link	sn_sc_ordered_item_link	Custom	String
owner	sn_sc_owner	Custom	String
no_delivery_time_v2	sn_sc_no_delivery_time_v2	Custom	String

ServiceNow field name	Index field name	Description	Data type
cost	sn_sc_cost	Custom	String
no_quantity_v2	sn_sc_no_quantity_v2	Custom	String
recurring_price	sn_sc_recurring_price	Custom	String
list_price	sn_sc_list_price	Custom	String
syst_tags	sn_sc_sys_tags	Custom	String
billable	sn_sc_billable	Custom	String
picture	sn_sc_picture	Custom	String
display_price_property	sn_sc_display_price_property	Custom	String
taxonomy_topic	sn_sc_taxonomy_topic	Custom	String
delivery_plain_script	sn_sc_delivery_plain_script	Custom	String
location	sn_sc_location	Custom	String

Attachments

Amazon Q supports crawling [ServiceNow Online attachments](#) and offers the following attachment field mappings.

ServiceNow field name	Index field name	Description	Data type
size_bytes	sn_file_size	Custom	Long (numeric)

ServiceNow field name	Index field name	Description	Data type
file_name	sn_file_name	Custom	String
sys_mod_count	sn_sys_mod_count	Custom	String
average_image_color	sn_average_image_color	Custom	String
image_width	sn_image_width	Custom	String
sys_updated_on	_last_updated_at	Default	Date
sys_tags	sn_sys_tags	Custom	String
table_name	sn_table_name	Custom	String
sys_id	sn_sys_id	Custom	String
image_height	sn_image_height	Custom	String
sys_updated_by	sn_updated_by	Custom	String
content_type	sn_content_type	Custom	String
sys_created_on	_created_at	Default	Date
size_compressed	sn_size_compressed	Custom	String
compressed	sn_compressed	Custom	String
state	sn_state	Custom	String
table_sys_id	sn_table_sys_id	Custom	String
chunk_size_bytes	sn_chunk_size_bytes	Custom	String
hash	sn_hash	Custom	String
sys_created_by	_authors	Default	String list

ServiceNow field name	Index field name	Description	Data type
sys_updated_by	sn_updated_by	Custom	String
url	sn_url	Custom	String
displayUrl	_source_uri	Default	String

Incidents

Amazon Q supports crawling [ServiceNow Online incidents](#) and offers the following incident field mappings.

ServiceNow field name	Index field name	Description	Data type
short_description	sn_inc_short_description	Custom	String
description	sn_inc_description	Custom	String
sys_updated_on	_last_updated_at	Default	Date
number	sn_inc_number	Custom	String
sys_updated_by	sn_updatedBy	Custom	String
displayUrl	_source_uri	Default	String
opened_by	sn_inc_opened_by	Custom	String
sys_created_on	_created_at	Default	Date
state	sn_inc_state	Custom	String
sys_created_by	_authors	Default	String list
business_impact	sn_inc_business_impact	Default	String

ServiceNow field name	Index field name	Description	Data type
impact	sn_inc_business_impact	Custom	String
priority	sn_inc_priority	Custom	String
urgency	sn_inc_urgency	Custom	String
opened_at	an_inc_opened_at	Custom	String
business_duration	sn_inc_business_duration	Custom	String
caller_id	sn_inc_caller_id	Custom	String
resolved_at	sn_inc_resolved_at	Custom	String
category	sn_inc_category	Custom	String
subcategory	sn_inc_subcategory	Custom	String
close_code	sn_inc_close_code	Custom	String
assignment_group	sn_inc_assignment_group	Custom	String
close_notes	sn_inc_close_notes	Custom	String
displayUrl	_source_uri	Default	String
sys_class_name	sn_inc_sys_class_name	Custom	String
parent_incident	an_inc_parent_incident	Custom	String
incident_state	sn_incident_state	Custom	String
company	sn_inc_company	Custom	String

ServiceNow field name	Index field name	Description	Data type
assigned_to	sn_inc_assigned_to	Custom	String
hold_reason	an_inc_hold_reason	Custom	String
work_notes	sn_inc_work_notes	Custom	String
comments_and_work_notes	sn_inc_comments_and_work_notes	Custom	String
work_notes_list	sn_work_notes_list	Custom	String
comments	sn_inc_comments	Custom	String
sys_id	sn_inc_sys_id	Custom	String
url	sn_url	Custom	String
active	sn_inc_active	Custom	String
activity_due	sn_inc_activity_due	Custom	String
additional_assignee_list	sn_inc_additional_assign_list	Custom	String
approval	sn_inc_approval	Custom	String
approval_history	sn_inc_approval_history	Custom	String
approval_set	sn_inc_approval_set	Custom	Date
business_service	sn_inc_business_service	Custom	String
closed_by	sn_inc_closed_by	Custom	String
cmdb_ci	sn_inc_cmdb_id	Custom	String

ServiceNow field name	Index field name	Description	Data type
resolved_by	sn_inc_resolved_by	Custom	String
sys_domain	sn_inc_sys_domain	Custom	String
business_stc	sn_inc_business_stc	Custom	String
calendar_duration	sn_inc_calendar_duration	Custom	String
calendar_stc	sn_inc_calendar_stc	Custom	String
cause	sn_inc_cause	Custom	String
caused_by	sn_inc_caused_by	Custom	String
child_incidents	sn_inc_child_incidents	Custom	String
closed_at	sn_inc_closed_at	Custom	String
contact_type	sn_inc_contact_type	Custom	String
contract	sn_inc_contract	Custom	String
correlation_display	sn_inc_correlation_display	Custom	String
delivery_plan	sn_inc_delivery_plan	Custom	String
delivery_task	sn_inc_delivery_task	Custom	String
due_date	sn_inc_due_date	Custom	String
escalation	sn_inc_escalation	Custom	String
expected_start	sn_inc_expected_start	Custom	String

ServiceNow field name	Index field name	Description	Data type
follow_up	sn_inc_follow_up	Custom	String
group_list	sn_inc_group_list	Custom	String
knowledge	sn_inc_knowledge	Custom	String
location	sn_inc_location	Custom	String
made_sla	sn_inc_made_sla	Custom	String
notify	sn_inc_notify	Custom	String
order	sn_inc_order	Custom	String
origin_id	sn_inc_origin_id	Custom	String
origin_table	sn_inc_origin_table	Custom	String
parent	sn_inc_parent	Custom	String
problem_id	sn_inc_problem_id	Custom	String
reassignment_count	sn_inc_reassignment_count	Custom	String
reopen_count	sn_inc_reopen_count	Custom	String
reopened_by	sn_inc_reopened_by	Custom	String
reopened_time	sn_inc_reopened_time	Custom	String
rfc	sn_inc_rfc	Custom	String
route_reason	sn_inc_route_reason	Custom	String
service_offering	sn_inc_service_offering	Custom	String

ServiceNow field name	Index field name	Description	Data type
severity	sn_inc_severity	Custom	String
sla_due	sn_inc_sla_due	Custom	Date
task_effective_number	sn_inc_task_effective_number	Custom	String
time_worked	sn_inc_time_worked	Custom	Date
universal_request	sn_inc_universal_request	Custom	String
upon_approval	sn_inc_upon_approval	Custom	String
upon_reject	sn_inc_upon_reject	Custom	String
user_input	sn_inc_user_input	Custom	String
watch_list	sn_inc_watch_list	Custom	String
work_end	sn_inc_work_end	Custom	String
work_start	sn_inc_work_start	Custom	String

IAM role for Amazon Q Business ServiceNow Online connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the BatchPutDocument and BatchDeleteDocument operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToGetSecret",
      "Effect": "Allow",
      "Action": [
        "secretsmanager:GetSecretValue"
      ],
      "Resource": [
        "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
      ]
    },
    {
      "Sid": "AllowsAmazonQToDecryptSecret",
      "Effect": "Allow",
      "Action": [
        "kms:Decrypt"
      ],
      "Resource": [
        "arn:aws:kms:{{region}}:{{account_id}}:key/[key_id]"
      ],
      "Condition": {
        "StringLike": {
          "kms:ViaService": [
            "secretsmanager.*.amazonaws.com"
          ]
        }
      }
    },
    {
      "Sid": "AllowsAmazonQToIngestDocuments",
      "Effect": "Allow",
```

```

    "Action": [
      "qbusiness:BatchPutDocument",
      "qbusiness:BatchDeleteDocument"
    ],
    "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
  },
  {
    "Sid": "AllowsAmazonQToIngestPrincipalMapping",
    "Effect": "Allow",
    "Action": [
      "qbusiness:PutGroup",
      "qbusiness:CreateUser",
      "qbusiness>DeleteGroup",
      "qbusiness:UpdateUser",
      "qbusiness>ListGroups"
    ],
    "Resource": [
      "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}",
      "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}",
      "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}/data-source/*"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateAndDeleteNI",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2>DeleteNetworkInterface"
    ],
    "Resource": [
      "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
      "arn:aws:ec2:{{region}}:{{account_id}}:security-group/[{{security_group}}]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterface",
      "ec2>DeleteNetworkInterface"
    ],
  },

```

```

"Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
"Condition": {
  "StringLike": {
    "aws:RequestTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
  },
  "ForAllValues:StringEquals": {
    "aws:TagKeys": [
      "AMAZON_Q"
    ]
  }
},
{
  "Sid": "AllowsAmazonQToCreateTags",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateTags"
  ],
  "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
  "Condition": {
    "StringEquals": {
      "ec2:CreateAction": "CreateNetworkInterface"
    }
  }
},
{
  "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateNetworkInterfacePermission"
  ],
  "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
  "Condition": {
    "StringLike": {
      "aws:ResourceTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
    }
  }
},
{
  "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
  "Effect": "Allow",
  "Action": [
    "ec2:DescribeNetworkInterfaces",
    "ec2:DescribeAvailabilityZones",

```

```

        "ec2:DescribeNetworkInterfaceAttribute",
        "ec2:DescribeVpcs",
        "ec2:DescribeRegions",
        "ec2:DescribeNetworkInterfacePermissions",
        "ec2:DescribeSubnets"
    ],
    "Resource": "*"
}
]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "{{source_account}}"
        },
        "ArnEquals": {
          "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
        }
      }
    }
  ]
}

```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Known limitations for the Amazon Q Business ServiceNow Online connector

The Amazon Q Business ServiceNow Online connector has the following known limitations:

- There is no REST API to wake up your ServiceNow Instance. You have to manually login into the ServiceNow instance to activate it when it's in hibernating mode.

Troubleshooting your Amazon Q Business ServiceNow Online connector

The following table provides information about error codes you may see for the ServiceNow Online connector and suggested troubleshooting actions.

Error code	Error message	Suggested resolution
SRN-5001	Error validating credentials due to invalid client id or client secret or username or password.	Provide a valid client id/client secret/username/password.
SRN-5002	Error validating credentials due to invalid username or password.	Provide a valid username/password.
SRN-5003	Access token is empty or null.	Provide a non empty or non null access token.
SRN-5004	Client ID exceeded the allowed length.	Provide a valid Client ID.
SRN-5005	Client Secret exceeded the allowed length.	Provide a valid Client Secret.
SRN-5006	Password exceeded the allowed length.	Provide a valid Password.
SRN-5007	clientSecret contains non-printable Ascii characters.	Provide a valid clientSecret.
SRN-5008	clientId contains non-printable Ascii characters.	Provide a valid clientId.

Error code	Error message	Suggested resolution
SRN-5009	servicenowInstance Version is not matching with hostUrl version.	Choose the correct servicenowInstance Version.
SRN-5010	Error validating credentials due to invalid username or password.	Provide a valid username/password.
SRN-5011	Amazon Q can't connect to the ServiceNow server with the specified credentials.	Provide admin credentials and try your request again.
SRN-5012	servicenowInstance Version is not matching with hostUrl version.	Choose the correct servicenowInstance Version (Tokyo/Sandiego/Rome.)
SRN-5013	ServiceNow instance is in hibernating mode.	Login to your ServiceNow instance before crawling.
SRN-5014	ServiceNow instance is not available.	Check your ServiceNow instance before crawling.
SRN-5100	Client id should not be empty.	Provide a valid client id.
SRN-5101	Client secret should not be empty.	Provide a valid client secret.
SRN-5102	User name should not be empty.	Provide a valid user name.
SRN-5103	Password should not be empty.	Provide a valid password.
SRN-5104	Auth type should not be empty.	Provide an auth Type.

Error code	Error message	Suggested resolution
SRN-5105	Incorrect auth type.	Auth type should be basicAuth or OAuth2.
SRN-5106	Host url should not be empty.	Provide a valid host url.
SRN-5120	crawlType should not be empty.	crawlType should be FORCED_FULL_CRAWL or FULL_CRAWL or CHANGE_LOG.
SRN-5122	isCrawlKnowledgeArticle should not be empty.	Provide valid isCrawlKnowledgeArticle.
SRN-5123	Invalid isCrawlKnowledgeArticle value.	isCrawlKnowledgeArticle should be true or false.
SRN-5124	isCrawlKnowledgeArticleAttachment should not be empty.	Provide valid isCrawlKnowledgeArticleAttachment.
SRN-5125	Invalid isCrawlKnowledgeArticleAttachment value.	isCrawlKnowledgeArticleAttachment should be true or false.
SRN-5126	isCrawlServiceCatalog should not be empty.	Provide valid isCrawlServiceCatalog.
SRN-5127	invalid isCrawlServiceCatalog value.	isCrawlServiceCatalog should be true or false.
SRN-5128	isCrawlServiceCatalogAttachment should not be empty.	Provide valid isCrawlServiceCatalogAttachment.
SRN-5129	Invalid isCrawlServiceCatalogAttachment value.	isCrawlServiceCatalogAttachment should be true or false.

Error code	Error message	Suggested resolution
SRN-5130	isCrawlIncident should not be empty.	Provide valid isCrawlIncident.
SRN-5131	invalid isCrawlIncident value.	isCrawlIncident should be true or false.
SRN-5132	isCrawlIncidentAttachment should not be empty.	Provide valid isCrawlIncidentAttachment.
SRN-5133	Invalid isCrawlIncidentAttachment value.	isCrawlIncidentAttachment should be true or false.
SRN-5134	Invalid incidentStateType.	Invalid incidentStateType. Incident State Type should be All, Open, Open - Unassigned or Resolved.
SRN-5135	applyACLForKnowledgeArticle should not be empty.	Provide valid applyACLForKnowledgeArticle.
SRN-5136	applyACLForServiceCatalog should not be empty.	Provide valid applyACLForServiceCatalog.
SRN-5137	applyACLForIncident should not be empty.	Provide valid applyACLForIncident.
SRN-5138	Invalid applyACLForKnowledgeArticle value.	applyACLForKnowledgeArticle should be true or false.
SRN-5139	Invalid applyACLForServiceCatalog value.	applyACLForServiceCatalog should be true or false.
SRN-5140	Invalid applyACLForIncident value.	applyACLForIncident should be true or false.

Error code	Error message	Suggested resolution
SRN-5141	invalid pattern :“file type pattern”	Provide valid patterns.
SRN-5142	includePublicArticlesOnly should not be empty.	Provide valid includePublicArticlesOnly.
SRN-5143	Invalid includePublicArticlesOnly value.	includePublicArticlesOnly should be true or false.
SRN-5144	Invalid URI.	Provide valid URI.
SRN-5145	isCrawlActiveServiceCatalog should not be empty.	Provide valid isCrawlActiveServiceCatalog.
SRN-5146	isCrawlInactiveServiceCatalog should not be empty.	Provide valid isCrawlInactiveServiceCatalog.
SRN-5147	isCrawlActiveIncident should not be empty.	Provide valid isCrawlActiveIncident.
SRN-5148	isCrawlInactiveIncident should not be empty.	Provide valid isCrawlInactiveIncident.
SRN-5149	Invalid isCrawlActiveServiceCatalog value.	isCrawlActiveServiceCatalog should be true or false.
SRN-5150	Invalid isCrawlInactiveServiceCatalog value.	isCrawlInactiveServiceCatalog should be true or false.
SRN-5151	Invalid isCrawlActiveIncident value.	isCrawlActiveIncident should be true or false.
SRN-5152	Invalid isCrawlInactiveIncident value.	isCrawlInactiveIncident should be true or false.

Error code	Error message	Suggested resolution
SRN-5153	servicenowInstance Version should not be empty.	Provide a valid servicenowInstance Version.
SRN-5154	The ServiceNow host name is invalid.	The ServiceNow host name should follow the format: example.service-now.com
SRN-5501	continuableInternalServerError.	Try again later.

Connecting Slack to Amazon Q Business

Slack is an enterprise communications app that lets users send messages and attachments through various public and private channels. You can connect your Slack instance to Amazon Q Business—using either the AWS Management Console, CLI, or the [CreateDataSource](#) API—and create an Amazon Q web experience.

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).
- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [Slack connector overview](#)
- [Prerequisites for connecting Amazon Q Business to Slack](#)
- [Setting up Slack for connecting to Amazon Q](#)
- [Connecting Amazon Q Business to Slack using the console](#)
- [Connecting Amazon Q Business to Slack using APIs](#)
- [How Amazon Q Business connector crawls Slack ACLs](#)

- [Amazon Q Business Slack data source connector field mappings](#)
- [IAM role for Amazon Q Business Slack connector](#)
- [Known limitations for the Amazon Q Business Slack connector](#)

Slack connector overview

The following table gives an overview of the Amazon Q Business Slack connector and its supported features.

Category	Feature	Support
Security	Authentication type	OAuth 2.0
	Authentication credentials	<ul style="list-style-type: none"> • Slack workspace ID • Either Slack Bot token or User token <p>User token lets you make API requests on behalf of the user. Bot token lets you make API requests as a Slack bot.</p>
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Identity crawling	Yes
	VPC	Yes
Crawl features	Custom metadata	No
	Entities	<p>Yes. The following entities are supported:</p> <ul style="list-style-type: none"> • Attachments (Files) • Text snippets • Posts • Text messages

Category	Feature	Support
		<ul style="list-style-type: none"> • Thread replies
	Field mappings	Yes. Supports default and custom field mappings. For more information, see Field mappings .
	Filters	<p>Yes. The following filters are supported :</p> <ul style="list-style-type: none"> • Crawl public channel • Crawl private channel • Crawl group messages • Crawl private messages • Crawl channel by type • Crawl channel by name • Including and excluding content by file type • Including and excluding content based on file name
	Sync mode	Supports full and incremental sync.
	File types	Supports all files supported by Amazon Q.
	Crawled as a document	<ul style="list-style-type: none"> • Each message • Each message attachment • Each channel post

Prerequisites for connecting Amazon Q Business to Slack

Before you begin, make sure that you have completed the following prerequisites.

In Slack, make sure you have:

- Created a Slack Bot User OAuth token or Slack User OAuth token. You can choose either token to connect Amazon Q to your Slack data source. See [Slack documentation on access tokens](#) for more information.

 **Note**

If you use the bot token as part of your Slack credentials, you cannot index direct messages and group messages. You must add the bot token to the channel you want to index.

- Noted your Slack workspace team ID from your Slack workspace main page URL. For example, <https://app.slack.com/client/T0123456789/...> where *T0123456789* is the team ID.
- Added the following OAuth scopes/ read permissions:

User token scope	Bot token scope
• channels:history	• channels:history
• channels:read	• channels:manage
• emoji:read	• channels:read
• files:read	• channels:read
• groups:history	• conversations.connect:manage
• groups:read	• conversations.connect:read
• im:history	• files:read
• im:read	• groups:history
• mpim:history	• groups:read
• mpim:read	• im:history
• team:read	• im:read
• users.profile:read	• mpim:history
• users:read	• mpim:read
• users:read.email	• reactions:read
	• team:read
	• usergroups:read
	• users.profile:read

User token scope**Bot token scope**

- users:read
- users:read.email

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your Slack authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Setting up Slack for connecting to Amazon Q

Before you connect Slack to Amazon Q, you need to create and retrieve the Slack credentials you will use to connect Slack to Amazon Q. You will also need to add any permissions needed by Slack to connect to Amazon Q.

The following procedure gives you an overview of how to configure Slack for connecting with Amazon Q.

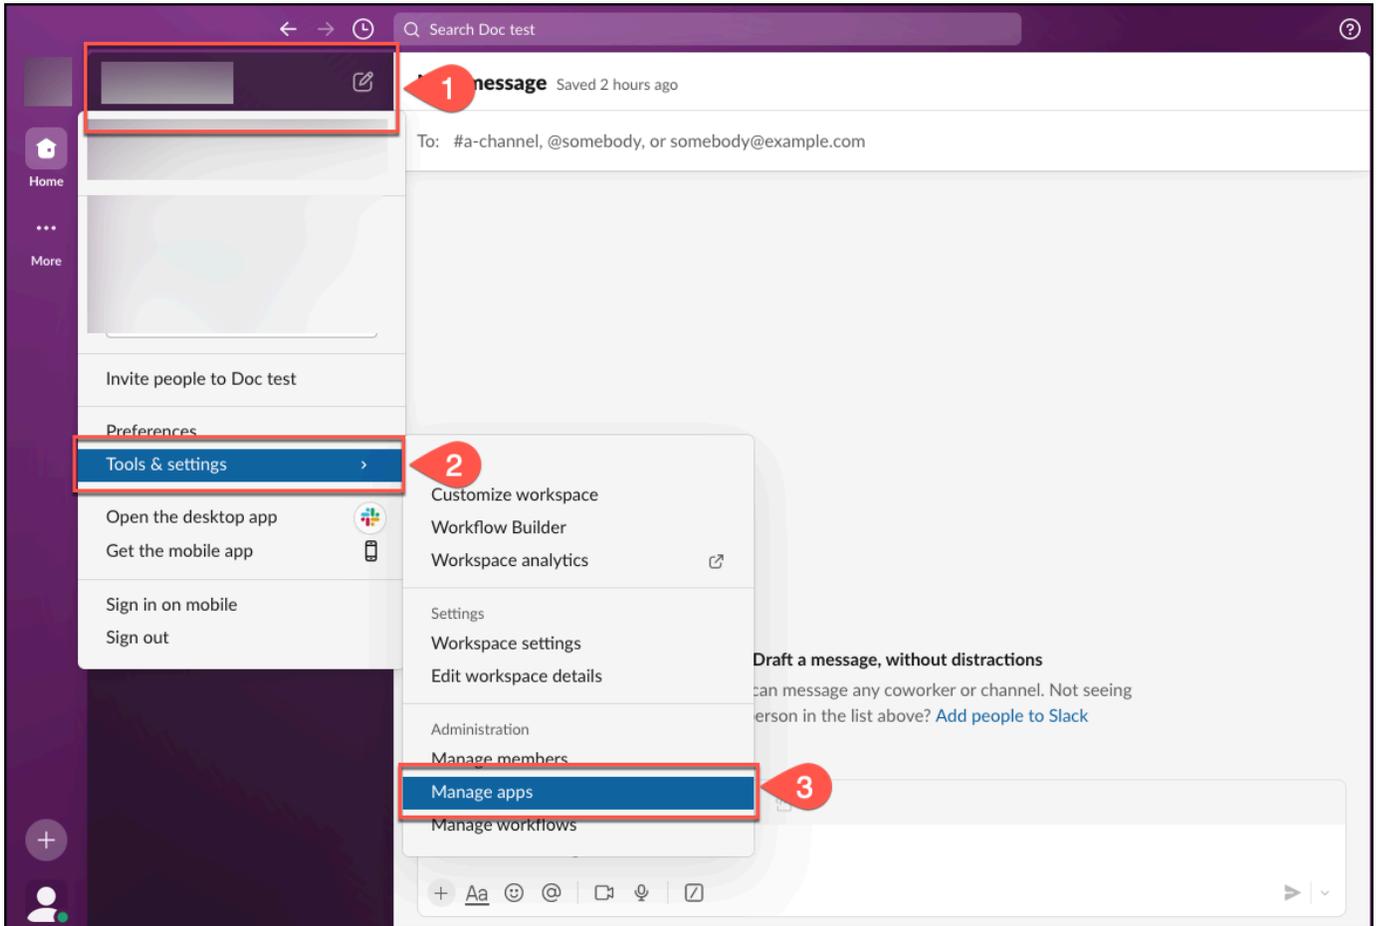
Configuring Slack authentication for Amazon Q

1. Log in to your Slack account and sign into your Slack workspace.

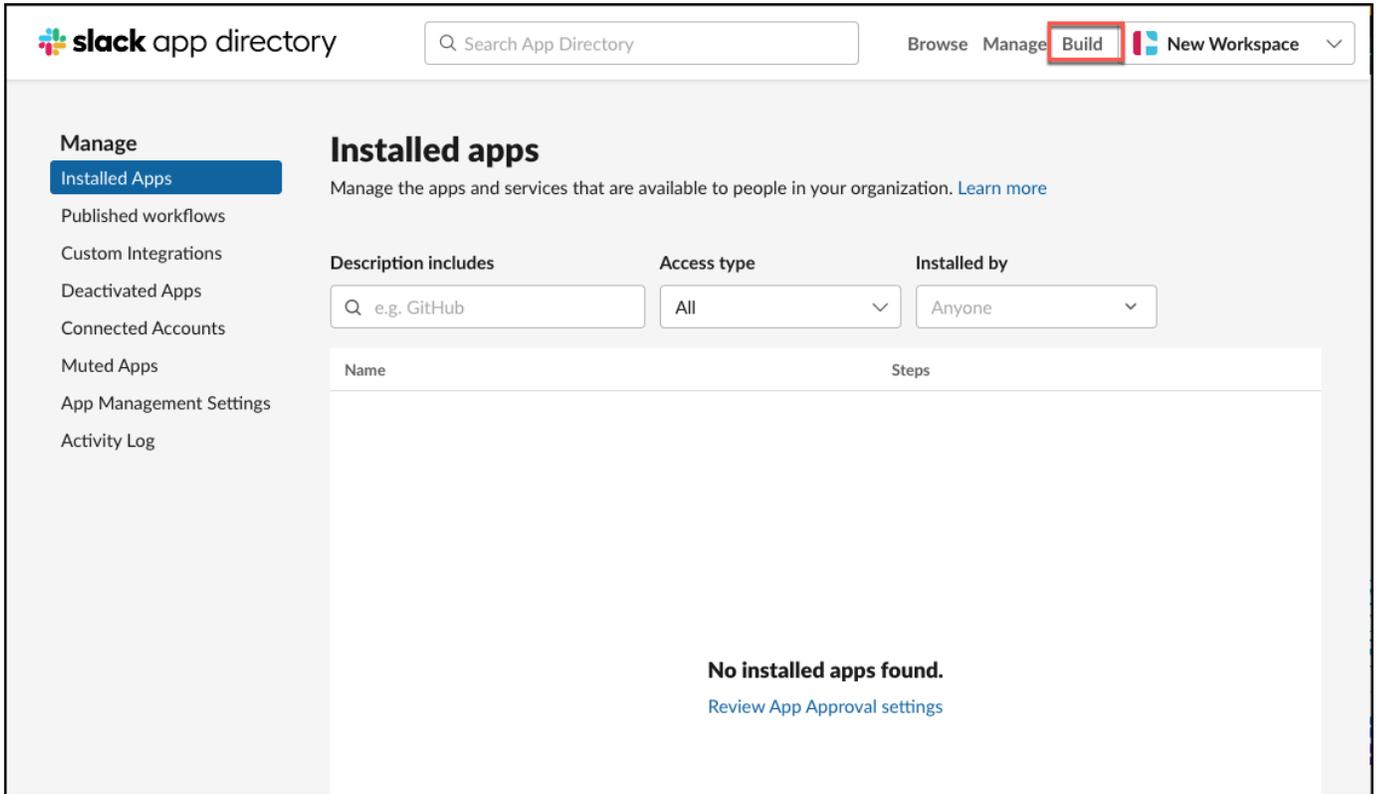
Note

To configure Slack for Amazon Q, you must be an admin user in the Slack account.

2. From the workspace menu, select **Tools and settings** and then select **Manage apps**.



3. From the **Slack App Directory** menu, select **Build**.



4. On the **Your Apps** page, select **Create an App**.

slack api

Q Search

Documentation Tutorials **Your Apps**

Automation >

Slack apps >

Messaging >

Surfaces >

Block Kit >

Enterprise >

Apps for Admins >

Gov Slack >

</> Reference

Your Apps

Build something amazing.

Use our APIs to build an app that makes people's working lives better. You can create an app that's just for your workspace or create a public Slack App to list in the App Directory, where anyone on Slack can discover it.

Create an App

Your App Configuration Tokens [Generate Token](#)

[Learn about tokens](#)

Don't see an app you're looking for? [Sign in to another workspace.](#)

Translated content

Stay updated

Changelog

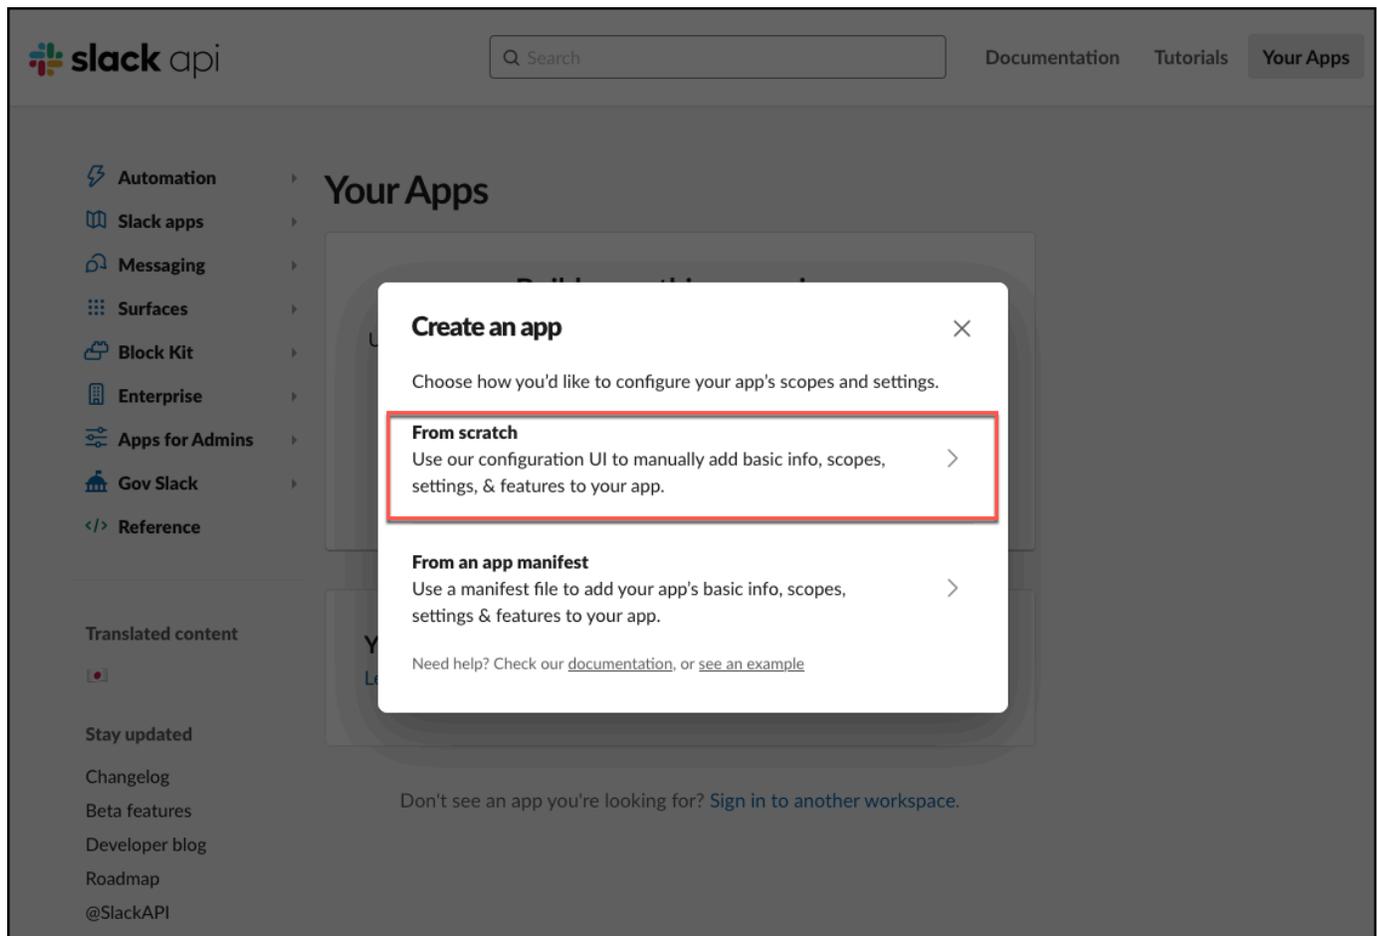
Beta features

Developer blog

Roadmap

@SlackAPI

5. On the **Create an app** page, select **From scratch**.



6. In the **Name app & choose workspace** dialog box that opens, add an **App name** and **Pick a workspace to deploy your app in**. Then select **Create App**.

Name app & choose workspace ×

App Name 1

e.g. Super Service

Don't worry - you'll be able to change this later.

Pick a workspace to develop your app in: 2

Select a workspace ▾

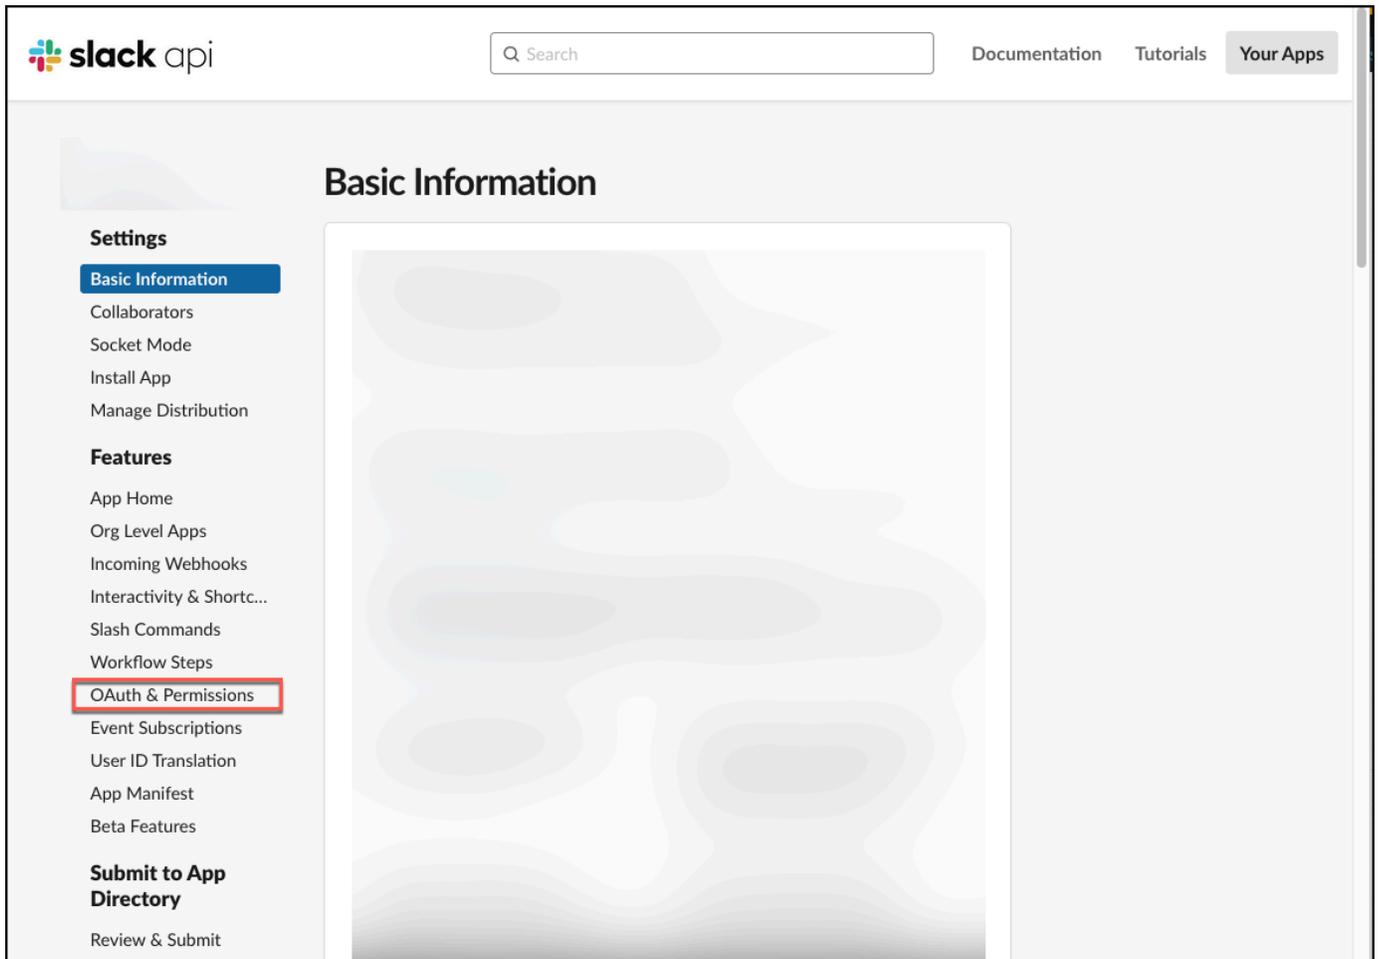
Keep in mind that you can't change this app's workspace later. If you leave the workspace, you won't be able to manage any apps you've built for it. The workspace will control the app even if you leave the workspace.

[Sign into a different workspace](#)

By creating a **Web API Application**, you agree to the [Slack API Terms of Service](#). 3

Cancel **Create App**

7. On the **Basic Information** page, from the **Settings** menu, select **OAuth & Permissions**.

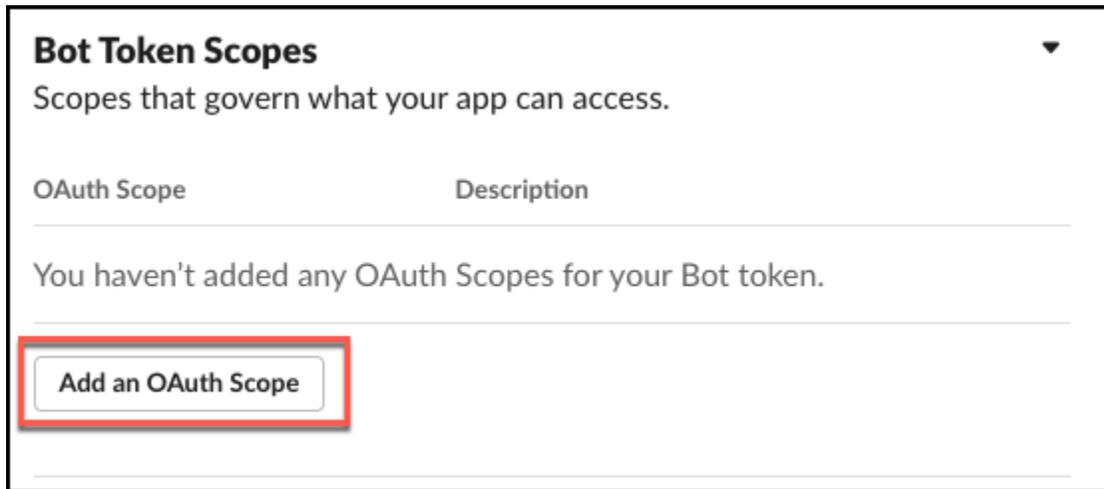


8. On the **OAuth & Permissions** page, go to **Scopes**, and then do the following based on whether you want to use a Bot Token to connect Slack to Amazon Q, or a User Token:

⚠ Important

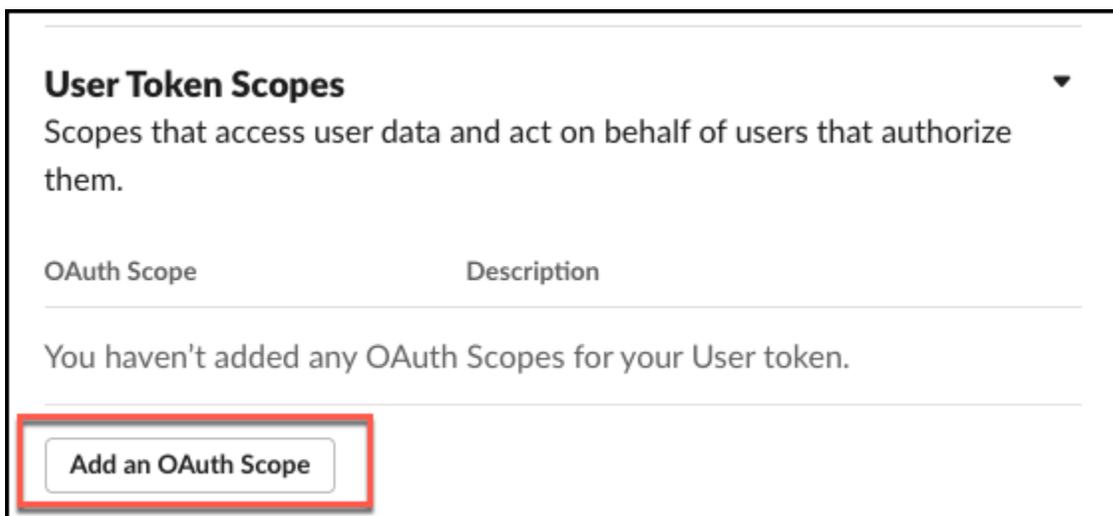
If you use the bot token as part of your Slack credentials, you cannot index direct messages and group messages, and you must add the bot token to the channel you want to index. For information on Slack token types, see [Token types](#) in Slack API.

- Add the following **Bot Token Scopes**:



- `channels:history` – View messages and other content in public channels that your app has been added to
- `channels:manage` – Manage public channels that your app has been added to and create new ones
- `channels:read` – View basic information about public channels in a workspace
- `conversations.connect:manage` – Receive Slack Connect invite events sent to the channels your app is in
- `conversations.connect:read` – Receive Slack Connect invite events sent to the channels your app is in
- `files:read` – View files shared in channels and conversations that your app has been added to
- `groups:history` – View messages and other content in private channels that your app has been added to
- `groups:read` – View basic information about private channels that your app has been added to
- `im:history` – View messages and other content in direct messages that your app has been added to
- `im:read` – View basic information about direct messages that your app has been added to
- `mpim:history` – View messages and other content in group direct messages that your app has been added to
- `mpim:read` – View basic information about group direct messages that your app has been added to

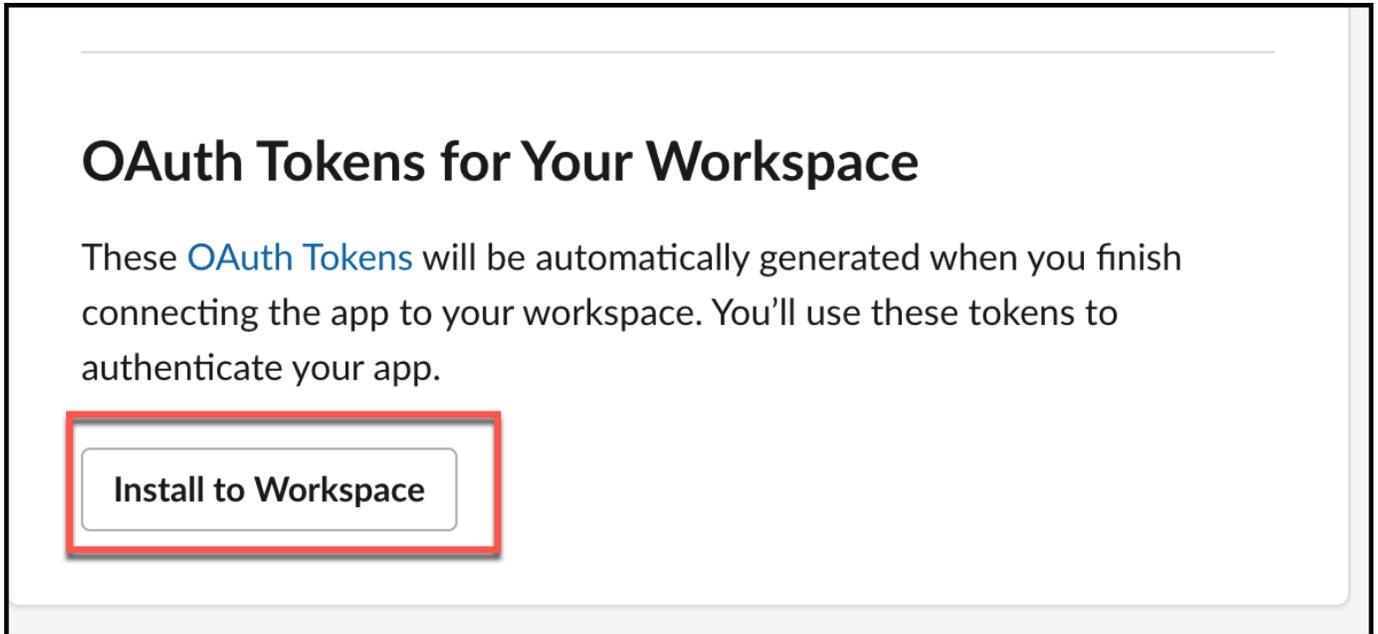
- `reactions:read` – View emoji reactions and their associated content in channels and conversations that your app has been added to
- `team:read` – View the name, email domain, and icon for workspaces your app is connected to
- `usergroups:read` – Create and manage user groups
- `users.profile:read` – View profile details about people in a workspace
- `users:read` – View people in a workspace
- `users:read.email` – View email addresses of people in a workspace
- Add the following **User Token Scopes**:



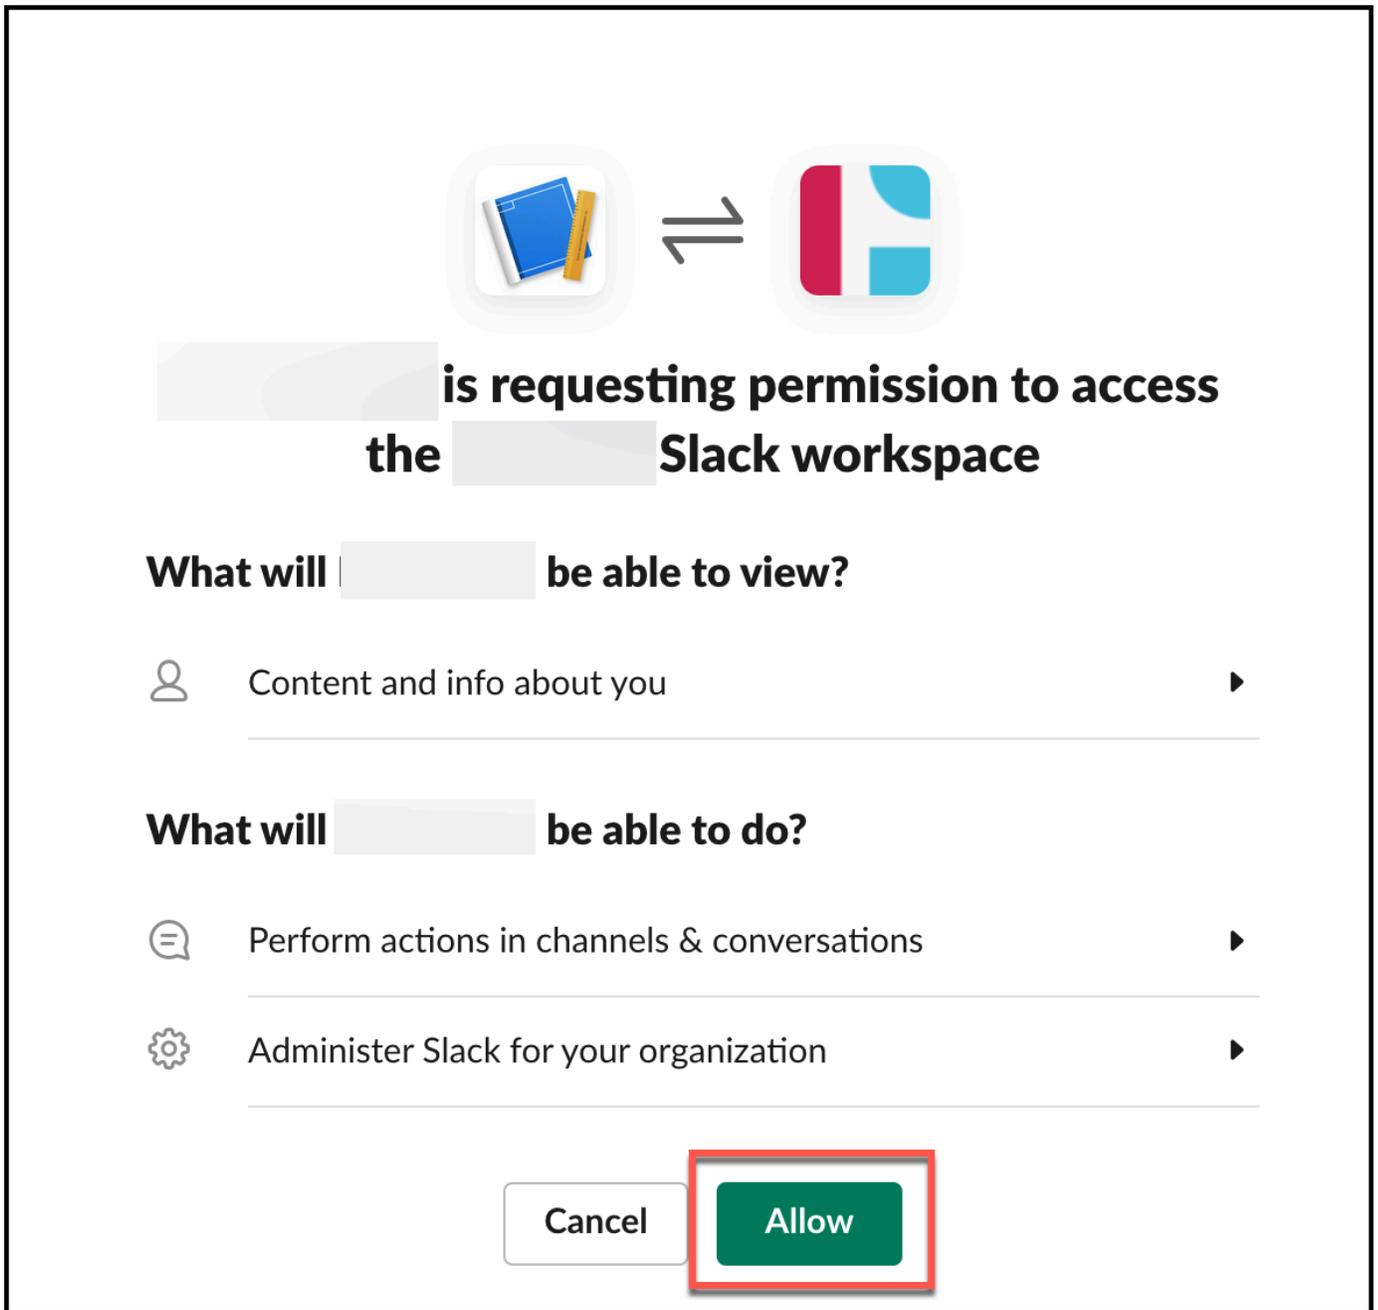
- `channels:history` – View messages and other content in a user's public channels
- `channels:read` – View basic information about public channels in a workspace
- `emoji:read` – View custom emoji in a workspace
- `files:read` – View files shared in channels and conversations that a user has access to
- `groups:history` – View messages and other content in a user's private channels
- `groups:read` – View basic information about a user's private channels
- `im:history` – View messages and other content in a user's direct messages
- `im:read` – View basic information about a user's direct messages
- `mpim:history` – View messages and other content in a user's group direct messages
- `mpim:read` – View basic information about a user's group direct messages
- `team:read` – View the name, email domain, and icon for workspaces a user is connected to

- `users.profile:read` – View profile details about people in a workspace
- `users.profile:read` – View profile details about people in a workspace
- `users:read` – View people in a workspace

9. Then, scroll to **OAuth Tokens for Your Workspace** section, and choose **Install to Workspace**.



10. On the dialog box that opens up informing you that the app that you created is requesting permission to access the Slack workspace you wanted to connect it to, select **Allow**.



On successful completion, the console will display a **OAuth Tokens for Your Workspace** screen.

11. From the **OAuth Tokens for Your Workspace** screen, copy and save the OAuth token you will use to connect to Amazon Q—either **User OAuth Token** or **Bot User OAuth Token**. You input this as **Slack token** when you connect to Amazon Q.

OAuth Tokens for Your Workspace

These tokens were automatically generated when you installed the app to your team. You can use these to authenticate your app. [Learn more.](#)

User OAuth Token

[Redacted token] Copy

Access Level: Workspace

Bot User OAuth Token

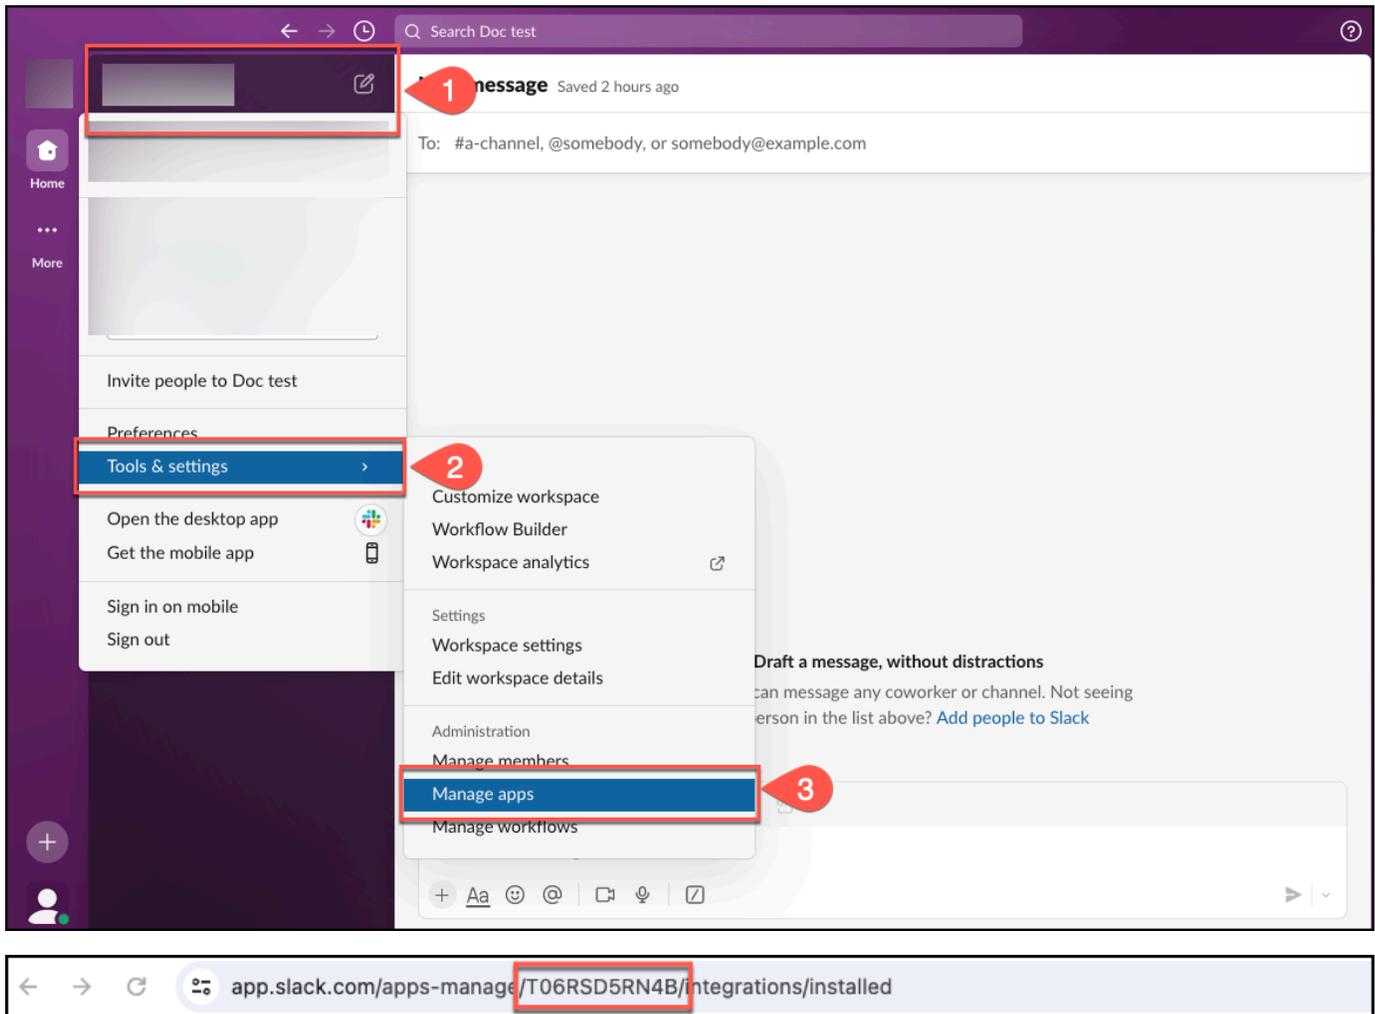
[Redacted token] Copy

Access Level: Workspace

[Reinstall to Workspace](#)

12. Next, you retrieve your Slack team ID. You need this to connect to Amazon Q.

From the Slack workspace menu, select **Tools and settings** and then select **Manage apps**. You'll find your team ID in the URL of the page that opens.



You now have the Slack Team ID and Slack token you need to connect to Amazon Q.

Connecting Amazon Q Business to Slack using the console

The following procedure outlines how to connect Amazon Q Business to Slack using the AWS Management Console.

Connecting Amazon Q to Slack

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).

4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **Slack** page, enter the following information:

6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. In **Source, Slack workspace team ID** – The team ID of your Slack workspace.
8. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

 **Note**

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

9. **Authentication** – Enter the following information for your **AWS Secrets Manager secret**.
 - a. **Secret name** – A name for your secret.
 - b. For **Slack token** – Enter the authentication credential values you created in your Slack account.
10. **Configure VPC and security group – optional** – Choose whether you want to use a VPC. If you do, enter the following information:
 - a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
 - b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

11. **Identity crawler** – Choose to activate Amazon Q identity crawler to sync identity information. For more information, see [Identity crawler](#).

12. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

13. In **Sync scope**, enter the following information:

- a. **Select type of content to crawl** – Select any combination of **All channels**, **Public channels**, **Private channels**, **Group messages**, and **Private messages**.
- b. **Select crawl start date** – Choose the date from which the Amazon Q connector will start crawling content.
- c. **Additional configuration – optional** – Configure the following settings:
 - In **Channels** (available only if you've chosen to crawl **Channels**), do the following:
 - **Channel ID/Name** – Choose between **Channel ID** and **Channel Name**.

 **Note**

You can choose to configure both.

- For **Channel ID** – Enter the **Channel ID**. The **Channel ID** filter applies to both public and private channels.
- For **Channel Name** – Choose the **Channel type** and enter the **Channel name**. You can select between **Public channel** and **Private channel**.

 **Note**

If you choose to configure filters for both **Channel ID** and **Channel Name**, the Amazon Q Slack connector will prioritize channel IDs over channel names. If you choose to configure filters for either **Channel ID** or **Channel Name**, the Amazon Q Slack connector will ignore **Private** and **Group** messages even if you've chosen to crawl private and group messages in **Sync scope**.

- In **Messages**, for **Select sync scope for content** – Choose to **Include bot messages**, and/or **Include archived messages**.
- **Regex patterns** – Add regex patterns to include or exclude file names or file types. You can add a total of 100 patterns. Examples of regex patterns include:
 - **File type** – .pdf, .docx

- **File name** – Hello*.txt, TestFile.*
14. For **Sync mode**, choose how you want to update your index when your data source content changes. When you sync your data source with Amazon Q for the first time, all content is synced by default.
 - **Full sync**—Sync all content regardless of the previous sync status.
 - **New, modified, or deleted content sync**—Sync only new, modified, and deleted documents.
 15. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
 16. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
 17. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:
 - a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
 - b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

 **Note**

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

18. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

19. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

Note

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to Slack using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the configuration parameter to provide a JSON schema with all other configuration information specific to your data source connector.

For an example of the API request, see [CreateDataSource](#) in the Amazon Q API Reference.

Slack JSON schema

The following is the Slack JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
      "type": "object",
      "properties": {
        "repositoryEndpointMetadata": {
          "type": "object",
          "properties": {
            "teamId": {
              "type": "string"
            }
          },
          "required": ["teamId"]
        }
      }
    },
    "repositoryConfigurations": {
      "type": "object",
      "properties": {
```

```

    "All": {
      "type": "object",
      "properties": {
        "fieldMappings": {
          "type": "array",
          "items": [
            {
              "type": "object",
              "properties": {
                "indexFieldName": {
                  "type": "string"
                },
                "indexFieldType": {
                  "type": "string",
                  "enum": ["STRING", "STRING_LIST", "DATE", "LONG"]
                },
                "dataSourceFieldName": {
                  "type": "string"
                },
                "dateFieldFormat": {
                  "type": "string",
                  "pattern": "yyyy-MM-dd'T'HH:mm:ss'Z'"
                }
              },
              "required": [
                "indexFieldName",
                "indexFieldType",
                "dataSourceFieldName"
              ]
            }
          ]
        },
        "required": [
          "fieldMappings"
        ]
      },
      "required": []
    },
    "additionalProperties": {
      "type": "object",
      "properties": {
        "isCrawlAcl": {

```

```
    "type": "boolean"
  },
  "fieldForUserId": {
    "type": "string"
  },
  "exclusionPatterns": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "inclusionPatterns": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "crawlBotMessages": {
    "type": "boolean"
  },
  "excludeArchived": {
    "type": "boolean"
  },
  "conversationType": {
    "type": "array",
    "items": {
      "type": "string",
      "enum": [
        "PUBLIC_CHANNEL",
        "PRIVATE_CHANNEL",
        "GROUP_MESSAGE",
        "DIRECT_MESSAGE"
      ]
    }
  },
  "channelFilter": {
    "type": "object",
    "properties": {
      "private_channel": {
        "type": "array",
        "items": {
          "type": "string"
        }
      }
    }
  },
},
```



```
    "public_channel": {
      "type": "array",
      "items": {
        "type": "string"
      }
    }
  },
  "channelIdFilter": {
    "type": "array",
    "items": {
      "type": "string"
    }
  },
  "sinceDate": {
    "anyOf": [
      {
        "type": "string",
        "pattern": "^[0-9]{4}-[0-9]{2}-[0-9]{2}T[0-9]{2}:[0-9]{2}:[0-9]{2}Z$"
      },
      {
        "type": "string",
        "pattern": ""
      }
    ]
  },
  "lookBack": {
    "type": "string",
    "pattern": "^[0-9]*$"
  }
},
"required": [
]
},
"syncMode": {
  "type": "string",
  "enum": [
    "FORCED_FULL_CRAWL",
    "FULL_CRAWL",
    "CHANGE_LOG"
  ]
},
"type" : {
  "type" : "string",
```

```

    "pattern": "SLACK"
  },
  "enableIdentityCrawler": {
    "type": "boolean"
  },
  "secretArn": {
    "type": "string"
  }
},
"version": {
  "type": "string",
  "anyOf": [
    {
      "pattern": "1.0.0"
    }
  ]
},
"required": [
  "connectionConfiguration",
  "repositoryConfigurations",
  "syncMode",
  "additionalProperties",
  "secretArn",
  "type",
  "enableIdentityCrawler"
]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint for the data source.
repositoryEndpointMetadata	The endpoint information for the data source.
teamId	The Slack team ID you copied from your Slack main page URL.

Configuration	Description
<code>repositoryConfigurations</code>	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings.
<ul style="list-style-type: none"> • All 	A list of objects that map the attributes or field names of your Slack pages and assets to Amazon Q index field names.
<code>additionalProperties</code>	Additional configuration options for your content in your data source.
<code>isCrawlAcl</code>	Specify <code>true</code> to crawl access control information from documents.
<code>fieldForUserId</code>	Specify field to use for <code>UserId</code> for ACL crawling.
<ul style="list-style-type: none"> • <code>inclusionPatterns</code> 	A list of regular expression patterns to include specific content in your Slack data source. Content that matches the patterns are included in the index. Content that doesn't match the patterns are excluded from the index. If any content matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence, and the content isn't included in the index.
<ul style="list-style-type: none"> • <code>exclusionPatterns</code> 	A list of regular expression patterns to exclude specific content in your Slack data source. Content that matches the patterns are excluded from the index. Content that doesn't match the patterns are included in the index. If any content matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence, and the content isn't included in the index.

Configuration	Description
<code>crawlBotMessages</code>	true to crawl Slack bot messages.
<code>excludeArchived</code>	true to exclude archived messages from crawl.
<code>conversationType</code>	The type of conversation that you want to index whether <code>PUBLIC_CHANNEL</code> , <code>PRIVATE_CHANNEL</code> , <code>GROUP_MESSAGE</code> and <code>DIRECT_MESSAGE</code> .
<code>channelFilter</code>	The type of channel that you want to index whether <code>private_channel</code> or <code>public_channel</code> .
<code>channelIdFilter</code>	You can choose to crawl specific channels by channel ID using the <code>channelIdFilter</code> .
<code>sinceDate</code>	You can choose to configure a <code>sinceDate</code> parameter so that the Slack connector crawls content based on a specific <code>sinceDate</code> .
<code>lookBack</code>	You can choose to configure a <code>lookBack</code> parameter so that the Slack connector crawls <code>lookBack</code> content.

Configuration	Description
syncMode	<p>Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose between the following options:</p> <ul style="list-style-type: none">• Use <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index.• Use <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index.• Use <code>CHANGE_LOG</code> to incrementally crawl only new and modified content each time your data source syncs with your index.
type	<p>The type of data source. Specify <code>SLACK</code> as your data source type.</p>
enableIdentityCrawler	<p>Specify <code>true</code> to use the Amazon Q identity crawler to sync identity/principal information on users and groups with access to specific documents.</p>
secretArn	<p>The Amazon Resource Name (ARN) of an AWS Secrets Manager secret that contains the key-value pairs required to connect to your Slack. The secret must contain a JSON structure with the following keys:</p> <pre data-bbox="829 1644 1507 1801">{ "slackToken": "<i>token</i>" }</pre>

Configuration	Description
version	The version of this template that's currently supported.

How Amazon Q Business connector crawls Slack ACLs

When you connect an Slack data source to Amazon Q Business, Amazon Q Business crawls ACL information attached to a document (user and group information) from your Slack instance. If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

The Slack user IDs are mapped as follows:

- `_user_id`—User IDs exist in Slack on messages and channels where there are set access permissions. They are mapped from the user emails as the IDs in Slack.

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business Slack data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q Business enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a attribute mapping already available, or if you want to map additional document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q Slack connector supports the following field mappings:

Slack field name	Index field name	Description	Data type
size	sl_gen_size	Custom	Long (numeric)
emojis	sl_gen_emojis	Custom	String list
title	al_gen_title	Custom	String
authors	al_gen_authors	Custom	String list
url	sl_gen_url	Custom	String
category	sl_gen_category	Custom	String
created_at	sl_gen_created_at	Custom	Date
last_updated_at	sl_gen_last_update d_at	Custom	String
msg_channel_id	sl_message_channel _id	Custom	String

Slack field name	Index field name	Description	Data type
msg_channel_name	sl_msg_channel_name	Custom	String

IAM role for Amazon Q Business Slack connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the `BatchPutDocument` and `BatchDeleteDocument` operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQToGetSecret",
      "Effect": "Allow",
      "Action": [
        "secretsmanager:GetSecretValue"
      ],
      "Resource": [
        "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
      ]
    }
  ],
}
```



```

{
  "Sid": "AllowsAmazonQToDecryptSecret",
  "Effect": "Allow",
  "Action": [
    "kms:Decrypt"
  ],
  "Resource": [
    "arn:aws:kms:{{region}}:{{account_id}}:key/[key_id]"
  ],
  "Condition": {
    "StringLike": {
      "kms:ViaService": [
        "secretsmanager.*.amazonaws.com"
      ]
    }
  }
},
{
  "Sid": "AllowsAmazonQToIngestDocuments",
  "Effect": "Allow",
  "Action": [
    "qbusiness:BatchPutDocument",
    "qbusiness:BatchDeleteDocument"
  ],
  "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
},
{
  "Sid": "AllowsAmazonQToIngestPrincipalMapping",
  "Effect": "Allow",
  "Action": [
    "qbusiness:PutGroup",
    "qbusiness:CreateUser",
    "qbusiness>DeleteGroup",
    "qbusiness:UpdateUser",
    "qbusiness:ListGroups"
  ],
  "Resource": [
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}",
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}",
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}/data-source/*"
  ]
}

```

```

},
{
  "Sid": "AllowsAmazonQToCreateAndDeleteNI",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateNetworkInterface",
    "ec2:DeleteNetworkInterface"
  ],
  "Resource": [
    "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[{{subnet_ids}}]",
    "arn:aws:ec2:{{region}}:{{account_id}}:security-group/[{{security_group}}]"
  ]
},
{
  "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateNetworkInterface",
    "ec2:DeleteNetworkInterface"
  ],
  "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
  "Condition": {
    "StringLike": {
      "aws:RequestTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
    },
    "ForAllValues:StringEquals": {
      "aws:TagKeys": [
        "AMAZON_Q"
      ]
    }
  }
},
{
  "Sid": "AllowsAmazonQToCreateTags",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateTags"
  ],
  "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
  "Condition": {
    "StringEquals": {
      "ec2:CreateAction": "CreateNetworkInterface"
    }
  }
}

```

```

    },
    {
      "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
      "Effect": "Allow",
      "Action": [
        "ec2:CreateNetworkInterfacePermission"
      ],
      "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
      "Condition": {
        "StringLike": {
          "aws:ResourceTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
        }
      }
    }
  ],
  {
    "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
    "Effect": "Allow",
    "Action": [
      "ec2:DescribeNetworkInterfaces",
      "ec2:DescribeAvailabilityZones",
      "ec2:DescribeNetworkInterfaceAttribute",
      "ec2:DescribeVpcs",
      "ec2:DescribeRegions",
      "ec2:DescribeNetworkInterfacePermissions",
      "ec2:DescribeSubnets"
    ],
    "Resource": "*"
  }
]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAmazonQServicePrincipal",
      "Effect": "Allow",
      "Principal": {
        "Service": "qbusiness.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
    }
  ]
}

```

```
    "Condition": {
      "StringEquals": {
        "aws:SourceAccount": "{{source_account}}"
      },
      "ArnEquals": {
        "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
      }
    }
  }
]
```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Known limitations for the Amazon Q Business Slack connector

The Amazon Q Business Slack connector has the following known limitations:

- Because of API limitations, Amazon Q Slack connector can only retrieve a maximum of 100 pages with 100 files per page from each channel. Given this, the Slack connector can only crawl a maximum number of 10000 files per channel.

Connecting Zendesk to Amazon Q Business

Zendesk is a customer relationship management system that helps businesses automate and enhance customer support interactions. You can connect a Zendesk instance to Amazon Q Business—using either the AWS Management Console or the [CreateDataSource](#) API—and create an Amazon Q web experience.

Learn more

- For an overview of the Amazon Q web experience creation process, see [Configuring an application](#).
- For an overview of connector features, see [Data source connector concepts](#).
- For information about connector configuration best practices, see [Connector configuration best practices](#).

Topics

- [Zendesk connector overview](#)
- [Prerequisites for connecting Amazon Q Business to Zendesk](#)
- [Setting up Zendesk for connecting to Amazon Q Business](#)
- [Connecting Amazon Q Business to Zendesk using the console](#)
- [Connecting Amazon Q Business to Zendesk using APIs](#)
- [How Amazon Q Business connector crawls Zendesk ACLs](#)
- [Amazon Q Business Zendesk data source connector field mappings](#)
- [IAM role for Amazon Q Business Zendesk connector](#)
- [Known limitations for the Amazon Q Business Zendesk connector](#)
- [Troubleshooting your Amazon Q Zendesk connector](#)

Zendesk connector overview

The following table gives an overview of the Amazon Q Business Zendesk connector and its supported features.

Category	Feature	Support
Security	Authentication type	OAuth 2.0
	Authentication credentials	<ul style="list-style-type: none"> • Zendesk Client ID • Zendesk Client secret • Zendesk username • Zendesk password
	Access Control List (ACL) crawling	Yes. For more information, see ACL crawling .
	Identity crawling	Yes
	VPC	Yes
Crawl features	Custom metadata	Yes
	Entities	Yes. The following entities are supported:

Category	Feature	Support
		<ul style="list-style-type: none">• Ticket• Ticket comment• Ticket comment attachment• Community topic• Community post• Community post comment• Article• Article attachment• Article comment <div data-bbox="829 779 1511 997"><p> Note Each instance of an entity is crawled as a single document.</p></div>
	Field mappings	Yes. Supports both default and custom field mappings. For more information, see Field mappings .

Category	Feature	Support
	Filters	<p>Yes. The following filters are supported:</p> <ul style="list-style-type: none"> • Organization name filter • Crawl tickets • Crawl ticket comments • Crawl ticket comment attachments • Crawl articles • Crawl article attachments • Crawl article comments • Crawl community topics • Crawl community posts • Crawl community post comments • Including and excluding content by file type • Including content based on a specific date
	<u>Sync mode</u>	Supports full and incremental sync.
	<u>File types</u>	Supports all files supported by Amazon Q.
	<u>Crawled as a document</u>	<ul style="list-style-type: none"> • Each ticket • Each ticket comment • Each ticket comment attachment • Each article • Each article attachment • Each article comment • Each community topic • Each community post • Each community post comment

Prerequisites for connecting Amazon Q Business to Zendesk

Before you begin, make sure that you have completed the following prerequisites.

In Zendesk, make sure you have:

- Created a Zendesk Suite (Professional/Enterprise) administrative account.
- Copied your Zendesk host URL. For example, `https://{sub-domain}.zendesk.com/`. You need this URL to allow Amazon Q to connect with your Zendesk data source.
- Generated Zendesk OAuth 2.0 credentials containing a client id, client secret, username, and password. You need these credentials to authenticate Amazon Q to access Zendesk.

In your AWS account, make sure you have:

- Created an [IAM role](#) for your data source and, if using the Amazon Q API, noted the ARN of the IAM role.
- Stored your Zendesk authentication credentials in an AWS Secrets Manager secret and, if using the Amazon Q API, noted the ARN of the secret.

Note

If you're a console user, you can create the IAM role and Secrets Manager secret as part of configuring your Amazon Q application on the console.

For a list of things to consider while configuring your data source, see [Data source connector configuration best practices](#).

Setting up Zendesk for connecting to Amazon Q Business

Before you connect Zendesk to Amazon Q Business, you need to create and retrieve the Zendesk credentials you will use to connect Zendesk to Amazon Q. You will also need to add any authorization permissions needed by Zendesk to connect to Amazon Q.

The following procedure gives you an overview of how to configure Zendesk for Amazon Q.

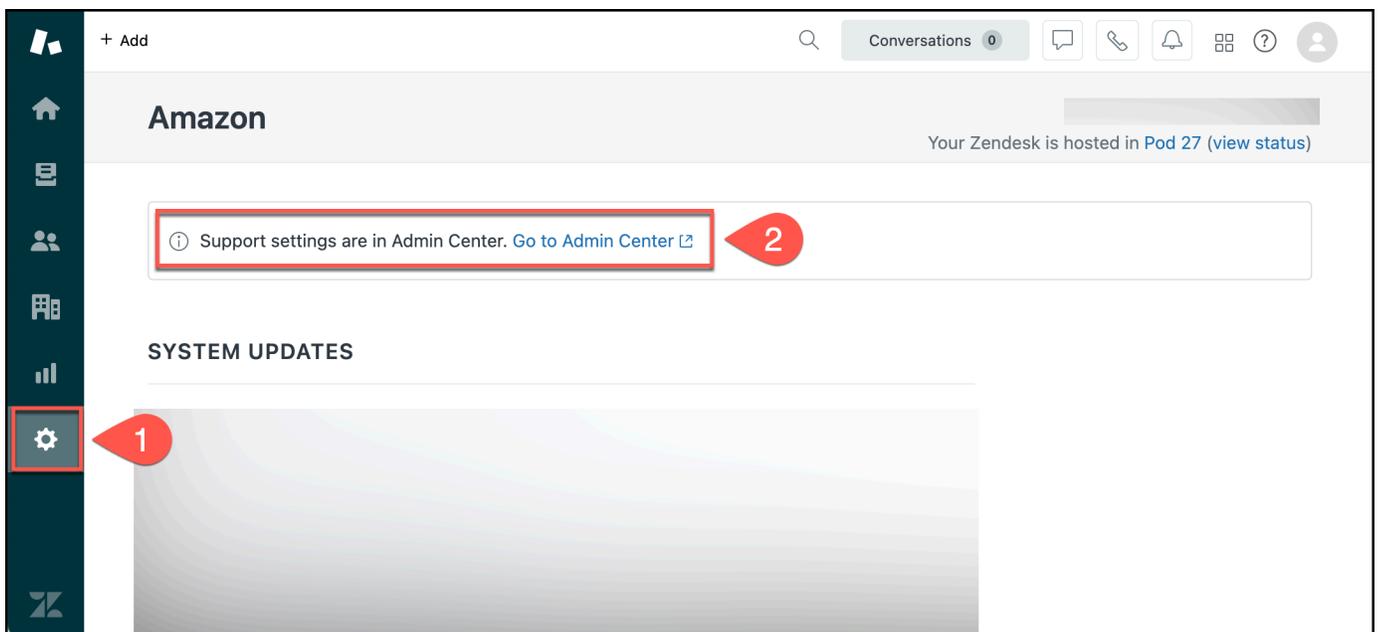
Configuring Zendesk for Amazon Q

1. Log in to your Zendesk account. Note the username and password you logged in with. You will need them later to connect to Amazon Q.
2. Copy your Zendesk URL, if you haven't already, from the Zendesk webpage URL. This will be the URL you will input as host URL in Amazon Q.

Note

You can also copy your Zendesk host URL from the top menu in the **Admin Center**.

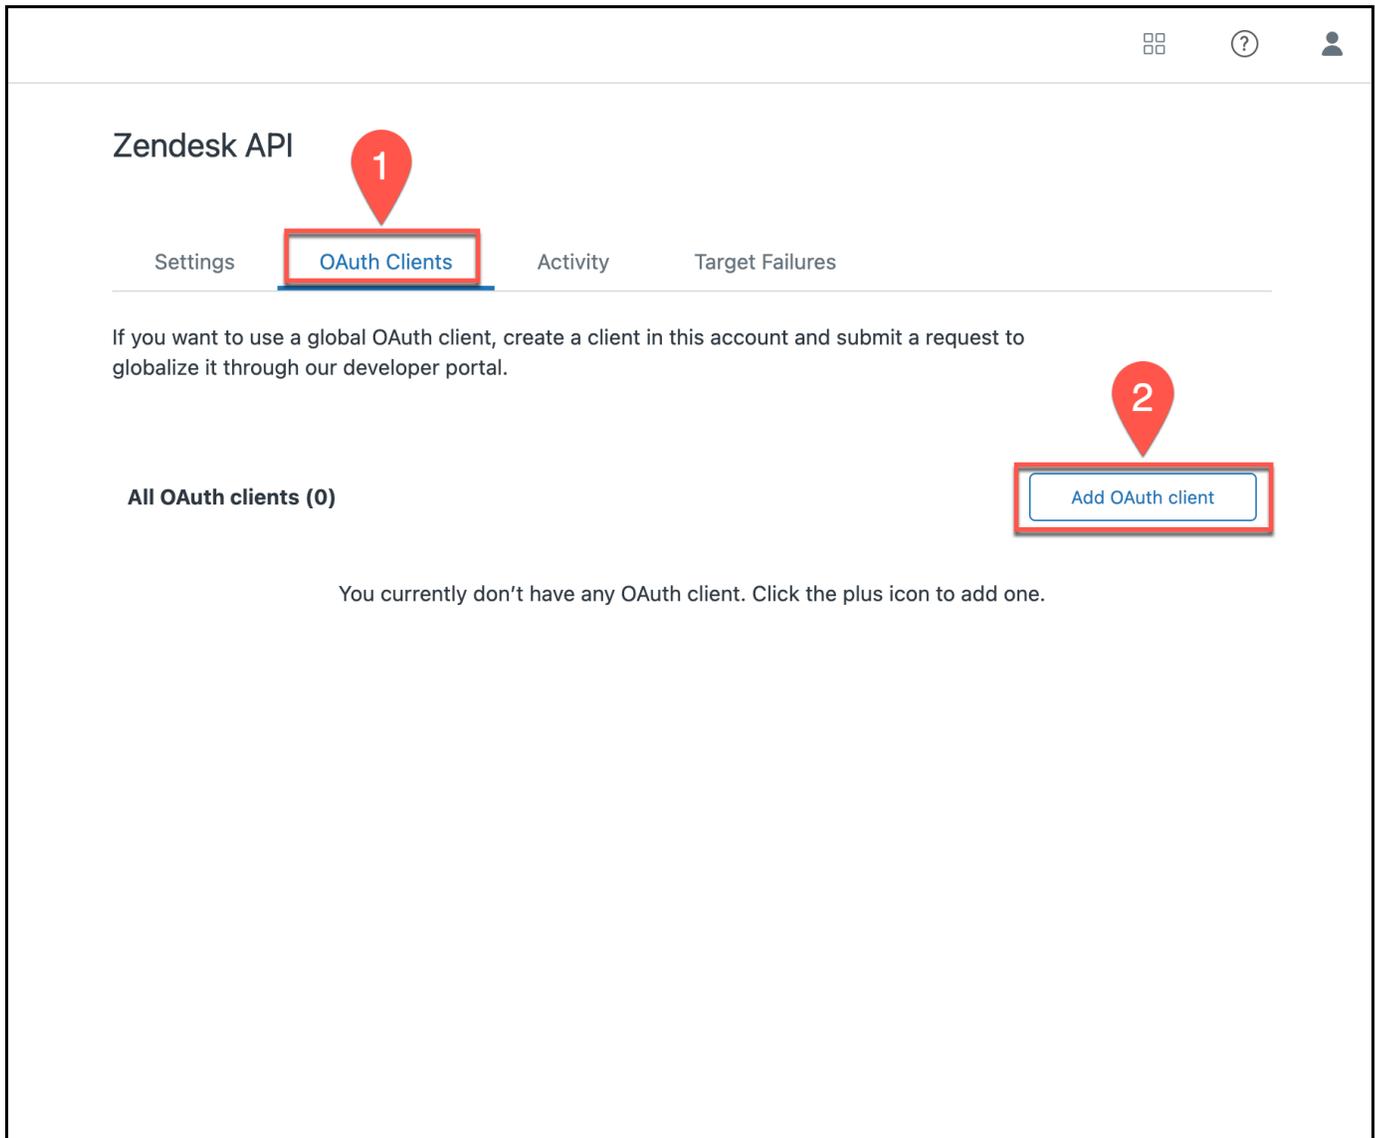
3. From the left navigation menu, choose the settings icon. Then, choose **Go to Admin Center**.



4. In **Admin Center**, from the left navigation menu, under **Apps and integrations**, choose **Zendesk API**.

The screenshot displays the Zendesk Admin Center interface. On the left sidebar, the 'Apps and integrations' menu item is highlighted with a red box and a red callout '1'. Below it, the 'APIs' menu item is also highlighted with a red box and a red callout '2'. The main content area shows the 'Admin Center' dashboard with several cards for 'Account', 'People', 'Channels', 'Workspaces', 'Objects and rules', and 'Help and support'.

5. From the **Zendesk API** menu, choose **OAuth Clients** and then choose **Add OAuth client**.



6. On the **OAuth Clients** page, under **Create a new OAuth client** enter the following information:
- **Client name** – A human-readable name for your client. This will be visible to users.
 - **Unique identifier** – An internal code-level identifier for your client. This will be the Client ID you input in Amazon Q.

Optionally, choose to fill in other information based on your use case. Then, choose **Save**.

Zendesk API

Settings

OAuth Clients

Activity

Target Failures

If you want to use a global OAuth client, create a client in this account and submit a request to globalize it through our developer portal.

Create a new OAuth client

1

Client name

Your client name shown to users when asked to grant access to your application or when viewing the list of apps that have been granted access.

Acme Integration for Zendesk

Description

A short description of your client for users when they're considering granting access to your application.

The Acme Integration for Zendesk allows your Acme account to connect securely to your Zendesk account to display Zendesk information in your Acme dashboard

Company

This name is displayed when users are asked to grant access to your application. The name helps users understand to whom they're granting access.

Amazon

Logo

Choose an image (JPG or PNG) to display when users are asked to grant access to your application.



2

Unique identifier

This is the name of your client for use in code. Example: my_awesome_app. This identifier is not shown to Zendesk users. You can change the initial suggestion. Identifiers with a zdg- prefix are reserved for global OAuth clients.

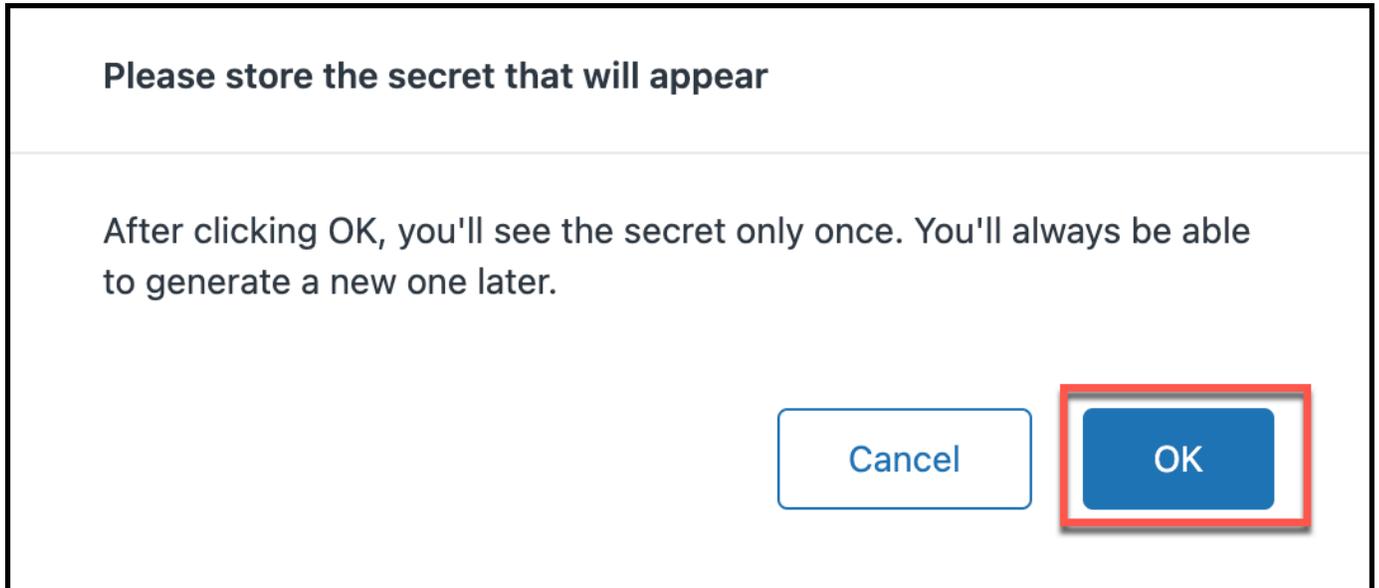
Redirect URLs

Specify the URL or URLs that Zendesk should use to redirect users after they decide whether or not to authorize your application to access Zendesk. The URLs must be absolute and not relative, https (unless localhost or 127.0.0.1), and newline-separated.

https://example.org/contact/oauth_redirect

3

7. On the **Please store the secret that will appear** dialog box that appears, select **OK**. Then, copy the secret you see into a text editor of your choice and save it. You won't be able to re-generate this secret so it's important that you store it securely. You will input this as the client secret during the connection configuration process in Amazon Q.



You now have the username, password, host URL, client ID, and client secret you need to connect Zendesk to Amazon Q.

Connecting Amazon Q Business to Zendesk using the console

The following procedure outlines how to connect Amazon Q Business to Zendesk using the AWS Management Console.

Connecting Amazon Q to Zendesk

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 5 data sources.

5. Then, on the **Zendesk** page, enter the following information:

6. **Name** – Name your data source for easy tracking.

Note: You can include hyphens (-) but not spaces. Maximum of 1,000 alphanumeric characters.

7. **Source** – Enter your **Zendesk URL**. For example, *https://{sub-domain (https://{host/})}.zendesk.com/*.
8. **Authorization** – Choose whether Amazon Q will crawl user and group access control list (ACL) information from your data source. Amazon Q can use this information to only generate responses from documents your end users have access to. See [Authorization](#) for more details.

Note

Using ACL data to filter responses is not a replacement for user authentication and authorization for your application. For information on setting up identity management for Amazon Q, see [Integrating with an Identity Provider \(IdP\)](#).

9. **Authentication** – Enter a name for your secret, a client ID, client secret, username, and password.
10. **Configure VPC and security group – optional** – Choose whether you want to use a VPC. If you do, enter the following information:
 - a. **Subnets** – Select up to 6 repository subnets that define the subnets and IP ranges the repository instance uses in the selected VPC.
 - b. **VPC security groups** – Choose up to 10 security groups that allow access to your data source. Ensure that the security group allows incoming traffic from Amazon EC2 instances and devices outside your VPC. For databases, security group instances are required.

For more information, see [VPC](#).

11. **Identity crawler** – Choose to activate Amazon Q identity crawler to sync identity information. For more information, see [Identity crawler](#).
12. **IAM role** – Choose an existing IAM role or create an IAM role to access your repository credentials and index content.

For more information, see [IAM role](#).

13. **Sync scope** – Set the content that you want to sync.
 - **Additional configuration – optional** – Configure the following settings:

- **Change log** – Select to update your index instead of syncing all your files.
 - **Organization name** – Enter the Zendesk organization names to filter your sync.
 - **Sync start date** – The date from which you want to index your content.
 - **Regex patterns** – Regular expression patterns to include or exclude certain files. You can add up to 100 patterns.
14. In **Sync run schedule**, for **Frequency** – Choose how often Amazon Q will sync with your data source. For more details, see [Sync run schedule](#).
15. **Tags - optional** – Add tags to search and filter your resources or track your AWS costs. See [Tags](#) for more details.
16. **Field mappings** – A list of data source document attributes to map to your index fields. Add the fields from the **Data source details** page after you finish adding your data source. You can choose from two types of fields:
- a. **Default** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can't edit these.
 - b. **Custom** – Automatically created by Amazon Q on your behalf based on common fields in your data source. You can edit these. You can also create and add new custom fields.

 **Note**

Support for adding custom fields varies by connector. You won't see the **Add field** option if your connector doesn't support adding custom fields.

For more information, see [Field mappings](#).

17. To finish connecting your data source to Amazon Q, select **Add data source**.

You are taken to the **Data source details**, where you can view your data source configuration details.

18. In **Data source details**, choose **Sync now** to allow Amazon Q to begin syncing (crawling and ingesting) data from your data source. When the sync job finishes, your data source is ready to use.

Note

You can also choose to view CloudWatch logs for your data source sync job by selecting **View CloudWatch logs**. If you get a Resource not found exception when you try to view your CloudWatch logs for a data source sync job in progress, it can be because the CloudWatch logs are not available yet. Wait for some time and check again.

Connecting Amazon Q Business to Zendesk using APIs

You use the [CreateDataSource](#) action to connect a data source to your Amazon Q application.

Then, you use the `configuration` parameter to provide a JSON schema with all other configuration information specific to your data source connector.

For an example of the API request, see [CreateDataSource](#) in the Amazon Q API Reference.

JSON schema

The following is the Zendesk JSON schema:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "connectionConfiguration": {
      "type": "object",
      "properties": {
        "repositoryEndpointMetadata": {
          "type": "object",
          "properties": {
            "hostUrl": {
              "type": "string",
              "pattern": "https:.*"
            }
          },
          "required": [
            "hostUrl"
          ]
        }
      },
      "required": [
        "hostUrl"
      ]
    }
  },
  "required": [
```



```

    "repositoryEndpointMetadata"
  ]
},
"repositoryConfigurations": {
  "type": "object",
  "properties": {
    "ticket": {
      "type": "object",
      "properties": {
        "fieldMappings": {
          "type": "array",
          "items": {
            "anyOf": [
              {
                "type": "object",
                "properties": {
                  "indexFieldName": {
                    "type": "string"
                  },
                  "indexFieldType": {
                    "type": "string",
                    "enum": ["STRING", "STRING_LIST", "LONG", "DATE"]
                  },
                  "dataSourceFieldName": {
                    "type": "string"
                  },
                  "dateFieldFormat": {
                    "type": "string",
                    "pattern": "dd-MM-yyyy HH:mm:ss"
                  }
                }
              },
              {
                "type": "string"
              }
            ]
          }
        },
        "required": [
          "indexFieldName",
          "indexFieldType",
          "dataSourceFieldName"
        ]
      }
    }
  }
},
"required": [
  "fieldMappings"
]

```

```

    ]
  },
  "ticketComment": {
    "type": "object",
    "properties": {
      "fieldMappings": {
        "type": "array",
        "items": {
          "anyOf": [
            {
              "type": "object",
              "properties": {
                "indexFieldName": {
                  "type": "string"
                },
                "indexFieldType": {
                  "type": "string",
                  "enum": ["STRING", "STRING_LIST", "LONG", "DATE"]
                },
                "dataSourceFieldName": {
                  "type": "string"
                },
                "dateFieldFormat": {
                  "type": "string",
                  "pattern": "dd-MM-yyyy HH:mm:ss"
                }
              }
            },
            {
              "type": "object",
              "properties": {
                "indexFieldName": {
                  "type": "string"
                },
                "indexFieldType": {
                  "type": "string",
                  "enum": ["STRING", "STRING_LIST", "LONG", "DATE"]
                },
                "dataSourceFieldName": {
                  "type": "string"
                }
              }
            }
          ]
        }
      },
      "required": [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
      ]
    }
  },
  "required": [
    "fieldMappings"
  ]
},
"ticketCommentAttachment": {
  "type": "object",

```

```

"properties": {
  "fieldMappings": {
    "type": "array",
    "items": {
      "anyOf": [
        {
          "type": "object",
          "properties": {
            "indexFieldName": {
              "type": "string"
            },
            "indexFieldType": {
              "type": "string",
              "enum": ["STRING", "STRING_LIST", "LONG", "DATE"]
            },
            "dataSourceFieldName": {
              "type": "string"
            },
            "dateFieldFormat": {
              "type": "string",
              "pattern": "dd-MM-yyyy HH:mm:ss"
            }
          }
        },
        "required": [
          "indexFieldName",
          "indexFieldType",
          "dataSourceFieldName"
        ]
      ]
    }
  },
  "required": [
    "fieldMappings"
  ]
},
"article": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": {
        "anyOf": [

```

```

        {
            "type": "object",
            "properties": {
                "indexFieldName": {
                    "type": "string"
                },
                "indexFieldType": {
                    "type": "string",
                    "enum": ["STRING", "STRING_LIST", "LONG", "DATE"]
                },
                "dataSourceFieldName": {
                    "type": "string"
                },
                "dateFieldFormat": {
                    "type": "string",
                    "pattern": "dd-MM-yyyy HH:mm:ss"
                }
            },
            "required": [
                "indexFieldName",
                "indexFieldType",
                "dataSourceFieldName"
            ]
        }
    ],
    "required": [
        "fieldMappings"
    ]
},
"communityPostComment": {
    "type": "object",
    "properties": {
        "fieldMappings": {
            "type": "array",
            "items": {
                "anyOf": [
                    {
                        "type": "object",
                        "properties": {
                            "indexFieldName": {
                                "type": "string"
                            }
                        }
                    }
                ]
            }
        }
    }
}

```

```

    },
    "indexFieldType": {
      "type": "string",
      "enum": ["STRING", "STRING_LIST", "LONG", "DATE"]
    },
    "dataSourceFieldName": {
      "type": "string"
    },
    "dateFieldFormat": {
      "type": "string",
      "pattern": "dd-MM-yyyy HH:mm:ss"
    }
  },
  "required": [
    "indexFieldName",
    "indexFieldType",
    "dataSourceFieldName"
  ]
}
]
}
},
"required": [
  "fieldMappings"
]
},
"articleComment": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": {
        "anyOf": [
          {
            "type": "object",
            "properties": {
              "indexFieldName": {
                "type": "string"
              },
              "indexFieldType": {
                "type": "string",
                "enum": ["STRING", "STRING_LIST", "LONG", "DATE"]
              }
            }
          }
        ]
      }
    }
  }
}

```

```

        "dataSourceFieldName": {
          "type": "string"
        },
        "dateFieldFormat": {
          "type": "string",
          "pattern": "dd-MM-yyyy HH:mm:ss"
        }
      },
      "required": [
        "indexFieldName",
        "indexFieldType",
        "dataSourceFieldName"
      ]
    }
  ]
}
},
"required": [
  "fieldMappings"
]
},
"articleAttachment": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": {
        "anyOf": [
          {
            "type": "object",
            "properties": {
              "indexFieldName": {
                "type": "string"
              },
              "indexFieldType": {
                "type": "string",
                "enum": ["STRING", "STRING_LIST", "LONG", "DATE"]
              },
              "dataSourceFieldName": {
                "type": "string"
              },
              "dateFieldFormat": {
                "type": "string",

```

```

        "pattern": "dd-MM-yyyy HH:mm:ss"
      }
    },
    "required": [
      "indexFieldName",
      "indexFieldType",
      "dataSourceFieldName"
    ]
  }
]
}
},
"required": [
  "fieldMappings"
]
},
"communityTopic": {
  "type": "object",
  "properties": {
    "fieldMappings": {
      "type": "array",
      "items": {
        "anyOf": [
          {
            "type": "object",
            "properties": {
              "indexFieldName": {
                "type": "string"
              },
              "indexFieldType": {
                "type": "string",
                "enum": ["STRING", "STRING_LIST", "LONG", "DATE"]
              },
              "dataSourceFieldName": {
                "type": "string"
              },
              "dateFieldFormat": {
                "type": "string",
                "pattern": "dd-MM-yyyy HH:mm:ss"
              }
            }
          },
          "required": [
            "indexFieldName",

```

```

        "indexFieldType",
        "dataSourceFieldName"
    ]
}
]
}
},
"required": [
    "fieldMappings"
]
}
},
"secretArn": {
    "type": "string",
    "minLength": 20,
    "maxLength": 2048
},
"additionalProperties": {
    "type": "object",
    "properties": {
        "isCrawlAcl": {
            "type": "boolean"
        },
        "fieldForUserId": {
            "type": "string"
        },
        "organizationNameFilter": {
            "type": "array"
        },
        "sinceDate": {
            "type": "string",
            "pattern": "^[0-9]{4}-[0-9]{2}-[0-9]{2} [0-9]{2}:[0-9]{2}:[0-9]{2}$"
        },
        "inclusionPatterns": {
            "type": "array"
        },
        "exclusionPatterns": {
            "type": "array"
        },
        "isCrawlTicket": {
            "type": "string"
        }
    }
},

```



```
    "isCrawTicketComment": {
      "type": "string"
    },
    "isCrawTicketCommentAttachment": {
      "type": "string"
    },
    "isCrawlArticle": {
      "type": "string"
    },
    "isCrawlArticleAttachment": {
      "type": "string"
    },
    "isCrawlArticleComment": {
      "type": "string"
    },
    "isCrawlCommunityTopic": {
      "type": "string"
    },
    "isCrawlCommunityPost": {
      "type": "string"
    },
    "isCrawlCommunityPostComment": {
      "type": "string"
    }
  }
},
"type": {
  "type": "string",
  "pattern": "ZENDESK"
},
"syncMode": {
  "type": "string",
  "enum": [
    "FULL_CRAWL",
    "FORCED_FULL_CRAWL",
    "CHANGE_LOG"
  ]
},
"enableIdentityCrawler": {
  "type": "boolean"
}
},
"version": {
  "type": "string",
```

```

    "anyOf": [
      {
        "pattern": "1.0.0"
      }
    ]
  },
  "additionalProperties": false,
  "required": [
    "connectionConfiguration",
    "repositoryConfigurations",
    "additionalProperties",
    "syncMode",
    "secretArn",
    "type"
  ]
}

```

The following table provides information about important JSON keys to configure.

Configuration	Description
connectionConfiguration	Configuration information for the endpoint of the data source.
repositoryEndpointMetadata	The endpoint information for the data source.
hostURL	The Zendesk host URL. For example, <i>https://yoursubdomain.zendesk.com</i> .
repositoryConfigurations	Configuration information for the content of the data source. For example, configuring specific types of content and field mappings.
<ul style="list-style-type: none"> • ticket • ticketComment • ticketCommentAttachment • article • articleComment • articleAttachment 	A list of Zendesk objects and their metadata attributes that Amazon Q crawls and maps to Amazon Q index field names. The Zendesk data source field names must exist in your Zendesk custom metadata.

Configuration	Description
<ul style="list-style-type: none"> • communityTopic • communityPost • communityPostComment 	
secretARN	The Amazon Resource Name (ARN) of an AWS Secrets Manager secret that contains the key-value pairs required to connect to your Zendesk. The secret must contain a JSON structure with the following keys: host URL, client ID, client secret, username, and password.
additionalProperties	Additional configuration options for your content in your data source.
isCrawlAcl	true to crawl Access Control Lists.
fieldForUserId	Specify field to use for UserId for ACL crawling.
organizationFilter	If you want, you can choose to index tickets that exist within a specific Organization
sinceDate	If you want, you can configure a sinceDate parameter so that the Zendesk connector will crawl based on the sinceDate .
inclusionPatterns	A list of regular expression patterns to <i>include</i> specific files in your Zendesk data source. Files that match the patterns are included in the index. Files that don't match the patterns are excluded from the index. If a file matches both an inclusion and exclusion pattern, the exclusion pattern takes precedence, and the file isn't included in the index.

Configuration	Description
exclusionPatterns	A list of regular expression patterns to <i>exclude</i> specific files in your Zendesk data source. Files that match the patterns are excluded from the index. Files that don't match the patterns are included in the index. If a file matches both an exclusion and inclusion pattern, the exclusion pattern takes precedence, and the file isn't included in the index.
<ul style="list-style-type: none">• isCrawlTicket• isCrawlTicketComment• isCrawlTicketCommentAttachment• isCrawlArticle• isCrawlArticleComment• isCrawlArticleAttachment• isCrawlCommunityTopic• isCrawlCommunityPost• isCrawlCommunityPostComment	Input true to index these types of content.
type	Specify ZENDESK as your data source type.

Configuration	Description
syncMode	<p>Specify whether Amazon Q should update your index by syncing all documents or only new, modified, and deleted documents. You can choose between the following options:</p> <ul style="list-style-type: none"> • Use <code>FORCED_FULL_CRAWL</code> to freshly re-crawl all content and replace existing content each time your data source syncs with your index. • Use <code>FULL_CRAWL</code> to incrementally crawl only new, modified, and deleted content each time your data source syncs with your index. • Use <code>CHANGE_LOG</code> to incrementally crawl only new and modified content each time your data source syncs with your index.
enableIdentityCrawler	<p>Specify <code>true</code> to activate identity crawler. Identity crawler is activated by default. Crawling identity information on users and groups with access to certain documents is useful for user context filtering. Search results are filtered based on the user or their group access to documents. See Identity crawler for more information.</p>
version	<p>The version of the template that's currently supported.</p>

How Amazon Q Business connector crawls Zendesk ACLs

When you connect an Zendesk data source to Amazon Q Business, Amazon Q Business crawls ACL information attached to a document (user and group information) from your Zendesk instance. If you choose to activate ACL crawling, the information can be used to filter chat responses to your end user's document access level.

The group and user IDs are mapped as follows:

- `_group_ids` – Group IDs exist in Zendesk tickets and articles where there are set access permissions. They are mapped from the names of the groups in Zendesk .
- `_user_id` – Group IDs exist in Zendesk tickets and articles where there are set access permissions. They are mapped from the user emails as the IDs in Zendesk .

For more information, see:

- [Authorization](#)
- [Identity crawler](#)
- [Understanding User Store](#)

Amazon Q Business Zendesk data source connector field mappings

To improve retrieved results and customize the end user chat experience, Amazon Q Business enables you to map document attributes from your data sources to fields in your Amazon Q index.

Amazon Q offers two kinds of attributes to map to index fields:

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data source to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

When you connect Amazon Q to a data source, Amazon Q automatically maps specific data source document attributes to fields within an Amazon Q index. If a document attribute in your data source doesn't have a attribute mapping already available, or if you want to map additional document attributes to index fields, use the custom field mappings to specify how a data source attribute maps to an Amazon Q index field. You create field mappings by editing your data source after your application and retriever are created.

To learn more about document attributes and how they work in Amazon Q, see [Document attributes and types in Amazon Q](#).

⚠ Important

During Preview, filtering using document attributes in chat is only supported through the API.

The Amazon Q Zendesk connector supports the following entities and the associated reserved and custom attributes.

Supported entities and field mappings

- [Tickets](#)
- [Ticket comments](#)
- [Ticket comment attachment](#)
- [Article](#)
- [Article comment](#)
- [Article comment attachment](#)
- [Community topic](#)
- [Community post](#)
- [Community post comment](#)

Tickets

Amazon Q supports crawling [Zendesk Tickets](#) and offers the following ticket field mappings.

Zendesk field name	Index field name	Description	Data type
ticketChannel	zd-channel	Custom	String
category	_category	Default	String
authors	_authors	Default	String list
assignee	zd_assignee	Custom	String
tags	zd_tags	Custom	String list

Zendesk field name	Index field name	Description	Data type
status	zd_status	Custom	String
sourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date
organizationName	zd_organization_name	Custom	String

Ticket comments

Amazon Q supports crawling [Zendesk Ticket Comments](#) and offers the following ticket comment field mappings.

Zendesk field name	Index field name	Description	Data type
category	_category	Default	String
authors	_authors	Default	String list
status	zd_status	Custom	String
sourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date
organizationName	zd_organization_name	Custom	String

Ticket comment attachment

Amazon Q supports crawling [Zendesk Ticket Comment Attachments](#) and offers the following ticket comment attachment field mappings.

Zendesk field name	Index field name	Description	Data type
category	_category	Default	String
authors	_authors	Default	String list
status	zd_status	Custom	String
sourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date
organizationName	zd_organization_name	Custom	String

Article

Amazon Q supports crawling [Zendesk Articles](#) and offers the following article field mappings.

Zendesk field name	Index field name	Description	Data type
authors	_authors	Default	String list
labels	zd_article_labels	Custom	String list
section	zd_article_section	Custom	String list
sourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date

Article comment

Amazon Q supports crawling [Zendesk Article Comments](#) and offers the following article comment field mappings.

Zendesk field name	Index field name	Description	Data type
authors	_authors	Default	String list
labels	zd_article_labels	Custom	String list
section	zd_article_section	Custom	String list
sourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date

Article comment attachment

Amazon Q supports crawling [Zendesk Article Comment Attachments](#) and offers the following article comment attachment field mappings.

Zendesk field name	Index field name	Description	Data type
authors	_authors	Default	String list
labels	zd_article_labels	Custom	String list
fileName	zd_file_name	Custom	String
fileType	_file_type	Default	String
fileSize	zd_file_size	Custom	Long (numeric)
section	zd_article_section	Custom	String list
sourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date

Community topic

Amazon Q supports crawling [Zendesk Community Topics](#) and offers the following community topic field mappings.

Zendesk field name	Index field name	Description	Data type
topicName	zd_topic_name	Custom	String
sourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date
category	_category	Default	String

Community post

Amazon Q supports crawling [Zendesk Community Posts](#) and offers the following community post field mappings.

Zendesk field name	Index field name	Description	Data type
postName	zd_post_name	Custom	String
topicName	zd_topic_name	Custom	String
sourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date
category	_category	Default	String

Community post comment

Amazon Q supports crawling [Zendesk Community Post Comments](#) and offers the following community post comment field mappings.

Zendesk field name	Index field name	Description	Data type
postName	zd_post_name	Custom	String
topicName	zd_topic_name	Custom	String
sourceUrl	_source_uri	Default	String
createdAt	_created_at	Default	Date
updatedAt	_last_updated_at	Default	Date
category	_category	Default	String

IAM role for Amazon Q Business Zendesk connector

If you use the AWS CLI or an AWS SDK, you must create an AWS Identity and Access Management (IAM) policy before you create an Amazon Q resource. When you call the operation, you provide the Amazon Resource Name (ARN) role with the policy attached.

If you use the AWS Management Console, you can create a new IAM role in the Amazon Q console or use an existing IAM role.

To connect your data source connector to Amazon Q, you must give Amazon Q an IAM role that has the following permissions:

- Permission to access the `BatchPutDocument` and `BatchDeleteDocument` operations to ingest documents.
- Permission to access the [User Store](#) API operations to ingest user and group access control information from documents.
- Permission to access your AWS Secrets Manager secret to authenticate your data source connector instance.
- **(Optional)** If you're using Amazon VPC, permission to access your Amazon VPC.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
```

```

    "Sid": "AllowsAmazonQToGetSecret",
    "Effect": "Allow",
    "Action": [
      "secretsmanager:GetSecretValue"
    ],
    "Resource": [
      "arn:aws:secretsmanager:{{region}}:{{account_id}}:secret:[secret_id]"
    ]
  },
  {
    "Sid": "AllowsAmazonQToDecryptSecret",
    "Effect": "Allow",
    "Action": [
      "kms:Decrypt"
    ],
    "Resource": [
      "arn:aws:kms:{{region}}:{{account_id}}:key/[key_id]"
    ],
    "Condition": {
      "StringLike": {
        "kms:ViaService": [
          "secretsmanager.*.amazonaws.com"
        ]
      }
    }
  }
},
{
  "Sid": "AllowsAmazonQToIngestDocuments",
  "Effect": "Allow",
  "Action": [
    "qbusiness:BatchPutDocument",
    "qbusiness:BatchDeleteDocument"
  ],
  "Resource": "arn:aws:qbusiness:{{region}}:{{source_account}}:application/
{{application_id}}/index/{{index_id}}"
},
{
  "Sid": "AllowsAmazonQToIngestPrincipalMapping",
  "Effect": "Allow",
  "Action": [
    "qbusiness:PutGroup",
    "qbusiness:CreateUser",
    "qbusiness>DeleteGroup",
    "qbusiness:UpdateUser",

```

```

    "qbusiness:ListGroupsWith",
  ],
  "Resource": [
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}",
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}",
    "arn:aws:qbusiness:{{region}}:{{account_id}}:application/{{application_id}}/
index/{{index_id}}/data-source/*"
  ]
},
{
  "Sid": "AllowsAmazonQToCreateAndDeleteNI",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateNetworkInterface",
    "ec2:DeleteNetworkInterface"
  ],
  "Resource": [
    "arn:aws:ec2:{{region}}:{{account_id}}:subnet/[subnet_ids]",
    "arn:aws:ec2:{{region}}:{{account_id}}:security-group/[security_group]"
  ]
},
{
  "Sid": "AllowsAmazonQToCreateAndDeleteNIForSpecificTag",
  "Effect": "Allow",
  "Action": [
    "ec2:CreateNetworkInterface",
    "ec2:DeleteNetworkInterface"
  ],
  "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
  "Condition": {
    "StringLike": {
      "aws:RequestTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
    },
    "ForAllValues:StringEquals": {
      "aws:TagKeys": [
        "AMAZON_Q"
      ]
    }
  }
},
{
  "Sid": "AllowsAmazonQToCreateTags",
  "Effect": "Allow",

```

```

    "Action": [
      "ec2:CreateTags"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringEquals": {
        "ec2:CreateAction": "CreateNetworkInterface"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToCreateNetworkInterfacePermission",
    "Effect": "Allow",
    "Action": [
      "ec2:CreateNetworkInterfacePermission"
    ],
    "Resource": "arn:aws:ec2:{{region}}:{{account_id}}:network-interface/*",
    "Condition": {
      "StringLike": {
        "aws:ResourceTag/AMAZON_Q": "qbusiness_{{account_id}}_{{application_id}}_*"
      }
    }
  },
  {
    "Sid": "AllowsAmazonQToDescribeResourcesForVPC",
    "Effect": "Allow",
    "Action": [
      "ec2:DescribeNetworkInterfaces",
      "ec2:DescribeAvailabilityZones",
      "ec2:DescribeNetworkInterfaceAttribute",
      "ec2:DescribeVpcs",
      "ec2:DescribeRegions",
      "ec2:DescribeNetworkInterfacePermissions",
      "ec2:DescribeSubnets"
    ],
    "Resource": "*"
  }
]
}

```

To allow Amazon Q to assume a role, you must also use the following trust policy:

```
{
```

```

"Version": "2012-10-17",
"Statement": [
  {
    "Sid": "AllowsAmazonQServicePrincipal",
    "Effect": "Allow",
    "Principal": {
      "Service": "qbusiness.amazonaws.com"
    },
    "Action": "sts:AssumeRole",
    "Condition": {
      "StringEquals": {
        "aws:SourceAccount": "{{source_account}}"
      },
      "ArnEquals": {
        "aws:SourceArn": "arn:aws:qbusiness:{{region}}:
{{source_account}}:application/{{application_id}}"
      }
    }
  }
]
}

```

For more information on Amazon Q data source connector IAM roles, see [IAM roles for Amazon Q data source connectors](#).

Known limitations for the Amazon Q Business Zendesk connector

The Amazon Q Business Zendesk connector has the following known limitations:

- Deleted and archived articles and their comments and attachments aren't supported in **Change log** mode since there are no SDK methods/REST API available for fetching deleted or archived articles.
- Archived articles aren't supported in **Full Crawl** mode since there are no SDK methods/REST API available for fetching archived articles.
- Deleted community topics, community posts, and their comments are not supported in **Change Log** mode since there are no SDK methods/REST API available for fetching deleted topics, deleted posts, and their comments
- The Zendesk connector can't fetch community topics (added, edited, or deleted), and community posts and their comments (added, edited, or deleted) based on timestamps in **Change log** mode.

Troubleshooting your Amazon Q Zendesk connector

The following table provides information about error codes you may see for the Zendesk connector and suggested troubleshooting actions.

Error code	Error message	Suggested resolution
ZND-5001	Error validating credentials due to invalid username or password.	Provide a valid username/password.
ZND-5002	Error validating credentials due to invalid client Id or client Secret.	Provide a valid Zendesk client Id or client Secret.
ZND-5100	The host URL is null or empty.	Provide a valid host Url.
ZND-5101	The username is null or empty.	Provide a valid username.
ZND-5102	The password is null or empty.	Provide a valid password.
ZND-5103	The Zendesk client Id is null or empty.	Provide a valid client Id.
ZND-5104	The Zendesk client Secret is null or empty.	Provide a valid client Secret.
ZND-5105	Invalid date format for field 'sinceDate'.	Date format should be yyyy-MM-dd HH:mm:ss.
ZND-5106	Invalid value for field 'sinceDate'.	Since date should not be greater than the current date.
ZND-5107	The datatype for the index field is invalid.	Only String, Date and Long formats are supported for field mappings.

Error code	Error message	Suggested resolution
ZND-5108	The isCrawTicket value is invalid.	isCrawTicket should be a boolean value true or false.
ZND-5109	The isCrawTicketComment value is invalid.	isCrawTicketComment should be a boolean value true or false.
ZND-5110	The isCrawTicketCommentAttachment value is invalid.	isCrawTicketCommentAttachment should be a boolean value true or false.
ZND-5111	The isCrawlArticle value is invalid.	isCrawlArticle should be a boolean value true or false.
ZND-5112	The isCrawlArticleComment value is invalid.	isCrawlArticleComment should be a boolean value true or false.
ZND-5113	The isCrawlArticleAttachment value is invalid.	isCrawlArticleAttachment should be a boolean value true or false.
ZND-5114	The isCrawlCommunityTopic value is invalid.	isCrawlCommunityTopic should be a boolean value true or false.
ZND-5115	The isCrawlCommunityPost value is invalid.	isCrawlCommunityPost should be a boolean value true or false.
ZND-5116	The isCrawlCommunityPostComment value is invalid.	isCrawlCommunityPostComment should be a boolean value true or false.
ZND-5117	Repository Configurations is null or empty.	Repository Configurations should not be null or empty value.

Error code	Error message	Suggested resolution
ZND-5118	The Host Url pattern is not valid.	Provide a valid host url. Ex: 'https://{sub-domain}.zendesk.com/' or 'https://{sub-domain}.zendesk.com'
ZND-5119	The URI is invalid.	Provide a valid URI.
ZND-5120	The personal access token is null or empty.	Provide a valid patToken.
ZND-5121	The auth type is incorrect .	The auth type should be OAuth2 or OAuth2-ImplicitGrantFlow.
ZND-5122	The accessToken provided is expired, revoked, malformed or invalid.	Provide valid accessToken.
ZND-5123	The access token doesn't have sufficient permission.	Check the user has sufficient permission to crawl.
ZND-5500	Unable to fetch data from Zendesk.	Check your Zendesk account plan/subscription: it may have expired.
ZND-5501	Unable to generate access token.	Check your Zendesk configuration and try again.
ZND-5502	There was an error parsing the field value. The size has exceeded the maximum allowable limit.	The maximum size permitted is 1000 characters for the fields.
ZND-5503	The url is invalid.	Provide valid URL.

Understanding Amazon Q Business User Store

With the Amazon Q Business *User Store* feature, end users see Amazon Q chat responses generated *only* from the documents that they have access to within an Amazon Q application. To achieve this, Amazon Q creates a mapping within the data sources attached to that application. The mapping is between every unique user accessing the application and all the user IDs and user groups that they are associated with. Amazon Q stores this principal mapping information in its internal User Store. During chat, Amazon Q uses the mapping information to return answers that are scoped to a user's identity.

When you use the API, you use the User Store API actions to customize and configure your user management solution. For more details, see [Using User Store APIs](#).

When you use the console, you can configure Amazon Q to automatically crawl user and group information using **Authorization** and **Identity crawling** features during the connector setup process. You can't create, add, or customize users and groups to the user store using the AWS Management Console.

Note

The User Store feature is not available for the Amazon S3 and Amazon Q Web Crawler connectors that are used with Amazon Q. For more information about using access control information for user identity specific chat responses for these connectors, see [Amazon S3](#) and [Amazon Q Web Crawler](#).

Topics

- [Principal mapping](#)
- [How the User Store works](#)

Principal mapping

Amazon Q Business uses *principal mapping* to map users and groups with permissions to access an Amazon Q application to their user ids and group membership information within the data sources that are connected to the application.

Although user and group mapping is a synchronous, simultaneous process, the following sections explain them separately for conceptual clarity.

Topics

- [User mapping](#)
- [Group mapping](#)

User mapping

Each Amazon Q Business application can have multiple data sources connected to it. Each data source can have specific users and groups configured within it. Additionally, a user can be associated with multiple groups within a data source, or be attached to multiple groups across multiple data sources. A user attached to multiple data sources can also have different user IDs within these data sources.

A unique end user who signs in to an Amazon Q application must see only chat responses generated from documents that they have access to. To achieve that objective, Amazon Q maps their user IDs and group IDs within each data source to their identity provider (IdP) login credentials. Then, Amazon Q creates a universally unique identifier (UUID) to assign to each user. Using the UUID that it creates, Amazon Q stores a comprehensive mapping of the user's group membership in an application. During chat, Amazon Q checks this UUID that's stored in its user store and retrieves user access information to generate chat responses.

The User Store feature also supports the following user management scenarios:

- **An end user leaves your organization.**

When an end user leaves your organization, you can choose to delete the user from your user store.

- **An end user leaves your organization, and their email gets recycled.**

Because User Store assigns each user a UUID for secure and accurate chat responses, email recycling doesn't impact the content that a user sees. Any new user within your application that's using a recycled email ID will be assigned a new UUID to be used for response generation.

- **An end user with multiple login IDs needs chat content generated from documents they access using both these login IDs.**

With User Store, you can store user aliases attached to end user UUIDs. For example, a username Saanvi Sarkar uses two login IDs to sign in to Amazon Q—`saanvi_sarkar` and `saanvi_s`. You can store both IDs under the same UUID to ensure their chat responses are generated from content that they access using both login IDs.

Group mapping

Each Amazon Q Business application can have multiple data sources attached to it. Each data source in an Amazon Q application can have multiple groups attached to it. Multiple groups can repeat across multiple data sources. Additionally, each group across data sources can also contain multiple subgroups. Each Amazon Q application also has an associated identity provider (IdP) that can contain group information for the users accessing the application.

A unique end user signing in to an Amazon Q application must see only chat responses generated from documents within groups that they have access to. To achieve that objective, Amazon Q does the following:

- Automatically crawls local groups and their associated relationships from data sources during the connector configuration process.
- Provides you with API operations to map your end users group and subgroup membership details within each data source to their IdP group membership.

Then, Amazon Q creates a unique user identifier (UUID) to assign to each user. Under the UUID, Amazon Q stores a comprehensive mapping of the user's group membership in an application. During chat, Amazon Q checks this UUID that's stored in its user store and quickly retrieves group access information to generate chat responses.

The User Store feature supports the following group management scenarios:

- **Your users mapped to all groups that they have access to within an Amazon Q application.**

Amazon Q crawls all groups that a user has access to in a data source and stores this information under a user's UUID.

- **Create a subgroup of users within your application.**

For example, for a group called `company_employees`, you might want to create a subgroup `summer_interns` and specify group level access for the subgroup. You might also want to group your interns into further subgroups like `product_interns` and `engineering_interns`.

- **Map your data source groups to your IdP groups.**

A unique end user signing in to an Amazon Q application must see only chat responses generated from documents within groups they have access to. To support that objective, you can

use Amazon Q to map your end users group membership details within each data source to their IdP group membership.

 **Note**

Amazon Q doesn't interact or crawl this information from your IdP automatically. To ingest the relationship between data source groups and IdP groups, use the Amazon Q API.

How the User Store works

The following overview describes how principal mapping works by using either the console or the Amazon Q Business API.

Topics

- [Using the console](#)
- [Using the API](#)

Using the console

Each document in any data source has access control list (ACL) information inherently attached to it as metadata. ACLs contain information about which users and groups have access to a document. When you configure an Amazon Q Business data source connector, you can configure your connector to crawl this ACL information by activating the **Authorization** toggle. If you choose to crawl ACL, the connector automatically extracts and maps document access information internally.

When you crawl this ACL information, Amazon Q stores it in its internal user store to assess which user IDs have access to a document. If you choose to not crawl ACL information, all documents are considered public.

Each data source also contains information about the users and groups which have access to it. During data source connector configuration, you can choose to crawl the information about users and groups attached to each data source by using the **Identity crawler** feature. If you choose to crawl user and group information, the connector automatically extracts and maps user and group information internally.

Amazon Q stores this crawled identity information in the user store and uses it to match and map user and group ids with their document access details. You can only use the **Identity crawler** feature if you also crawl ACLs using the **Authorization** feature.

If you use the console, you must re-sync your data to your index to capture any changes in the ACL and user and group membership within your data source.

Using the API

When you configure your Amazon Q Business application, you use the following API operations to create your principal mapping solution:

User management

- [CreateUser](#) – Creates a universally unique identifier (UUID) that's mapped to a list of local user IDs within a data source.
- [DeleteUser](#) — Deletes a UUID that's mapped to a user.
- [UpdateUser](#) – Updates local user IDs within a data source that are mapped to a UUID.
- [GetUser](#) – Lists information associated with a user ID.

Group management

- [PutGroup](#) – Creates, or updates, a mapping of users to groups, or groups to subgroups. You can use this API operation to:
 - Map a group from groups in the data source to groups in your IdP.
 - Map a list of users and sub groups (for example, Interns) to a group (for example, Interns 2023).
- [DeleteGroup](#) – Deletes a group or a subgroup.
- [GetGroup](#) – Lists information about a group.

Using Amazon VPC with Amazon Q Business connectors

Amazon Q Business can connect to a virtual private cloud (VPC) that you created with Amazon Virtual Private Cloud to index content stored in data sources running in your private cloud. When you create a data source connector, you can provide security group and subnet identifiers for the subnet that contains your data source. With this information, Amazon Q creates an elastic network interface that it uses to securely communicate with your data source within your VPC.

To set up an Amazon Q data source connector with Amazon VPC, you can use either the AWS Management Console or the [CreateDataSource](#) API operation. If you use the console, you connect a VPC during the connector configuration process.

Note

The Amazon VPC feature is optional when setting up an Amazon Q data source connector. If your data source is accessible from the public internet, you don't need to enable the Amazon VPC feature. Not all Amazon Q data source connectors support Amazon VPC.

Important

During Preview, Amazon Q doesn't support AWS PrivateLink.

If your data source isn't running on Amazon VPC and isn't accessible from the public internet, you first connect your data source to your VPC using a virtual private network (VPN). Then, you can connect your data source to Amazon Q by using a combination of Amazon VPC and AWS Virtual Private Network. For information about setting up a VPN, see the [AWS VPN documentation](#).

Topics

- [Configuring Amazon VPC support for Amazon Q Business connectors](#)
- [Set up an Amazon Q Business data source to connect to Amazon VPC](#)
- [Using Amazon VPC with an Amazon S3 data source](#)
- [Connecting to a database in a VPC](#)
- [Troubleshooting VPC connection issues](#)

Configuring Amazon VPC support for Amazon Q Business connectors

To configure Amazon VPC for use with your Amazon Q Business connectors, take the following steps.

Steps

- [Step 1. Create Amazon VPC subnets for Amazon Q Business](#)
- [Step 2. Create Amazon VPC security groups for Amazon Q Business](#)

- [Step 3. Configure your external data source and Amazon VPC](#)

Step 1. Create Amazon VPC subnets for Amazon Q Business

Create or choose an existing Amazon VPC subnet that Amazon Q Business can use to access your data source. The prepared subnets must be in one of the following AWS Regions and Availability Zones:

- US West (Oregon)/us-west-2—usw2-az1, usw2-az2, usw2-az3
- US East (N. Virginia)/us-east-1—use1-az1, use1-az2, use1-az4

Your data source must be accessible from the subnets that you provided to Amazon Q connector.

For more information about how to configure Amazon VPC subnets, see [Subnets for your Amazon VPC](#) in the *Amazon VPC User Guide*.

If Amazon Q must route the connection between two or more subnets, you can prepare multiple subnets. For example, the subnet that contains your data source is out of IP addresses. In that case, you can provide Amazon Q with an additional subnet that has sufficient IP addresses and connected to the first subnet. If you list multiple subnets, the subnets must be able to communicate with each other.

Step 2. Create Amazon VPC security groups for Amazon Q Business

To connect your Amazon Q Business data source connector to Amazon VPC, you must prepare one or more security groups from your VPC to assign to Amazon Q. The security groups will be associated to the elastic network interface created by Amazon Q. This network interface controls inbound and outbound traffic to and from Amazon Q when accessing the Amazon VPC subnets.

Make sure that your security group's outbound rules allow the traffic from Amazon Q data source connectors to access the subnets and the data source that you are going to sync with. For example, you might use an MySQL connector to sync from a MySQL database. If you're using the default port, the security groups must allow Amazon Q to access port 3306 on the host that runs the database.

We recommend that you configure a default security group with the following values for Amazon Q to use:

- **Inbound rules** – If you choose to leave this empty, all inbound traffic will be blocked.

- **Outbound rules** – Add one rule to allow all outbound traffic so that Amazon Q can initiate the requests to sync from your data source.
 - **IP version** – IPv4
 - **Type** – All traffic
 - **Protocol** – All traffic
 - **Port range** – All
 - **Destination** – 0.0.0.0/0

For more information about how to configure Amazon VPC security groups, see [Security group rules](#) in the *Amazon VPC User Guide*.

Step 3. Configure your external data source and Amazon VPC

Make sure that your external data source has the correct permissions configuration and network settings for Amazon Q to access it. You can find detailed instructions on how to configure your data sources in the prerequisites section of each connector page.

Also, check your Amazon VPC settings and make sure that your external data source is reachable from the subnet you will assign to Amazon Q. To do this, we recommend that you create an Amazon EC2 instance in the same subnet with the same security groups and test access to your data source from this Amazon EC2 instance. For more information, see [Troubleshooting Amazon VPC connection](#).

Set up an Amazon Q Business data source to connect to Amazon VPC

When you add a new data source in Amazon Q Business, you can use the Amazon VPC feature if the selected data source connector supports this feature.

You can set up a new Amazon Q data source with Amazon VPC enabled by using the AWS Management Console or the Amazon Q API. Specifically, use the [CreateDataSource](#) API operation, and then use the `VpcConfiguration` parameter to provide the following information:

- `SubnetIds` – A list of identifiers of Amazon VPC subnets
- `SecurityGroupIds` – A list of identifiers of Amazon VPC security groups

If you use the console, you provide the required Amazon VPC information during connector configuration. To use the console to enable the Amazon VPC feature for a connector, you first

choose an Amazon VPC. Then, you provide identifiers of any Amazon VPC subnets and identifiers of any Amazon VPC security groups. You can choose the Amazon VPC subnets and Amazon VPC security groups that you created in [Configuring Amazon VPC](#), or use any existing ones.

Topics

- [Viewing Amazon VPC identifiers](#)
- [Checking your data source IAM role](#)

Viewing Amazon VPC identifiers

The identifiers for subnets and security groups are configured in the Amazon VPC console. To view the identifiers, use the following procedures.

To view subnet identifiers

1. Sign in to the AWS Management Console and open the Amazon VPC console at <https://console.aws.amazon.com/vpc/>.
2. From the navigation pane, choose **Subnets**.
3. From the **Subnets** list, choose the subnet that contains your database server.
4. From the **Details** tab, make a note of the identifier in the **Subnet ID** field.

To view security group identifiers

1. Sign in to the AWS Management Console and open the Amazon VPC console at <https://console.aws.amazon.com/vpc/>.
2. From the navigation pane, choose **Security groups**.
3. From the security group list, choose the group that you want the identifier for.
4. From the **Details** tab, make a note of the identifier in the **Security Group ID** field.

Checking your data source IAM role

Make sure that your data source connector AWS Identity and Access Management (IAM) role contains permissions to access your Amazon VPC.

If you use the console to create a new role for your IAM role, Amazon Q automatically adds the correct permissions to your IAM role on your behalf. If you use the API, or use an existing IAM role,

check that your role contains permissions to access Amazon VPC. To verify that you have the right permissions, see [IAM roles for data sources](#).

You can modify an existing data source to use a different Amazon VPC subnet. However, check your data source's IAM role and, if necessary, modify it to reflect the change for the Amazon Q data source connector to work properly.

Using Amazon VPC with an Amazon S3 data source

This topic provides a step-by-step example that shows how to connect to an Amazon S3 bucket by using an Amazon S3 connector through Amazon VPC. The example assumes that you're starting with an existing S3 bucket. We recommend that you upload just a few documents to your S3 bucket to test the example.

You can connect Amazon Q Business to your Amazon S3 bucket through Amazon VPC. To do so, you must specify the Amazon VPC subnet and Amazon VPC security groups when creating your Amazon S3 data source connector.

Important

So that an Amazon Q Amazon S3 connector can access your Amazon S3 bucket, make sure that you have assigned an Amazon S3 endpoint to your virtual private cloud (VPC). For more information about configuring an Amazon Q Amazon S3 connector with Amazon VPC, see [Using Amazon VPC with Amazon S3](#).

For Amazon Q to sync documents from your Amazon S3 bucket through Amazon VPC, you must complete the following steps:

- Set up an Amazon S3 endpoint for Amazon VPC. For more information about how to set up an Amazon S3 endpoint, see [Gateway endpoints for Amazon S3](#) in the *AWS PrivateLink Guide*.
- (Optional) Checked your Amazon S3 bucket policies to make sure that the Amazon S3 bucket is accessible from the virtual private cloud (VPC) that you assigned to Amazon Q. For more information, see [Controlling access from VPC endpoints with bucket policies](#) in the *Amazon S3 User Guide*.

Steps

- [Step 1: Configure an Amazon VPC](#)

- [\(Optional\) Step 2: Configure Amazon S3 bucket policy](#)
- [Step 3: Create a test Amazon S3 data source connector](#)

Step 1: Configure an Amazon VPC

Create a VPC network including a private subnet with an Amazon S3 gateway endpoint and a security group for Amazon Q to use later.

To configure a VPC with a private subnet, an S3 endpoint, and a security group

1. Sign in to the AWS Management Console and open the Amazon VPC console at <https://console.aws.amazon.com/vpc/>.
2. **Create a VPC with a private subnet and an S3 endpoint for Amazon Q to use:**

From the navigation pane, choose **Your VPCs**, and then choose **Create VPC**.

- a. For **Resources to create**, choose **VPC and more**.
- b. For **Name tag**, enable **Auto-generate**, then enter **qbusiness-s3-example**.
- c. For **IPv4 / IPv6 CIDR block**, keep the default values.
- d. For **Number of Availability Zones (AZs)**, choose **number 1**.
- e. Select **Customize AZs**, and then select an Availability Zone from the **First availability zone** list.

Amazon Q only supports a specific set of Availability Zones.

- f. For **Number of public subnets**, choose **number 0**.
- g. For **Number of private subnets**, choose **number 1**.
- h. For **NAT gateways**, choose **None**.
- i. For **VPC endpoints**, choose **Amazon S3 gateway**.
- j. Leave the rest of the values at their default settings.
- k. Select **Create VPC**.

Wait until the **Create VPC** workflow finishes. Then, choose **View VPC** to check the **VPC** you just created.

You have now created a VPC network with a private subnet, which does not have access to the public internet.

3. Copy your VPC endpoint ID of your Amazon S3 endpoint:

- a. From the navigation pane, choose **Endpoints**.
- b. In the **Endpoints** list, find the Amazon S3 endpoint `qbusiness-s3-example-vpce-s3` that you just created together with your VPC.
- c. Make a note of the **VPC endpoint ID**.

You have now created an Amazon S3 gateway endpoint to access your Amazon S3 bucket through a subnet.

4. Create a Security Group for Amazon Q to use:

- a. From the navigation pane, choose **Security Groups**, then select **Create security group**.
- b. For **Security group name**, enter `s3-data-source-security-group`.
- c. Choose your VPC from the **Amazon VPC** list.
- d. Leave **inbound rules** and **outbound rules** as the default.
- e. Choose **Create security group**.

You have now created a VPC security group.

You assign the subnet and security group that you created to your Amazon Q Amazon S3 data source connector during the connector configuration process.

(Optional) Step 2: Configure Amazon S3 bucket policy

In this optional step, learn how to configure an Amazon S3 bucket policy so that your Amazon S3 bucket is only accessible from the VPC that you assign to Amazon Q Business.

Amazon Q uses IAM roles to access your Amazon S3 bucket and doesn't require that you configure an Amazon S3 bucket policy. However, you might find it useful to create a bucket policy if you want to configure an Amazon S3 connector using an Amazon S3 bucket that has existing policies restricting access to it from the public internet.

To configure your Amazon S3 bucket policy

1. Open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. From the navigation pane, choose **Buckets**.

3. Choose the name of the Amazon S3 bucket that you want to sync with Amazon Q.
4. Choose the **Permissions** tab, scroll down to **Bucket policy**, and then click on **Edit**.
5. Add or modify your bucket policy to allow access only from the VPC endpoint that you created.

The following is an example bucket policy. Replace *bucket-name* and *vpce-id* with your Amazon S3 bucket name and the Amazon S3 endpoint ID that you noted earlier.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "Principal": "*",
      "Action": "s3:*",
      "Resource": "arn:aws:s3:::bucket-name/*",
      "Condition": {
        "StringNotEquals": {
          "aws:SourceVpce": "vpce-id"
        }
      }
    }
  ]
}
```

6. Select **Save changes**.

Your S3 bucket is now accessible only from the specific VPC that you created.

Step 3: Create a test Amazon S3 data source connector

To test your Amazon VPC configuration, create an Amazon S3 connector. Then, configure it with the VPC that you created by following the steps outlined in [Amazon S3](#).

For Amazon VPC configuration values, choose the values that you created during this example:

- **Amazon VPC(VPC)** – `qbusiness-s3-example-vpc`
- **Subnets** – `qbusiness-s3-example-subnet-private1-[availability zone]`
- **Security groups** – `s3-data-source-security-group`

Wait for your connector to finish creating. After the Amazon S3 connector has been created, choose **Sync now** to initiate a sync.

It might take several minutes to several hours to finish the sync, depending on how many documents are in your Amazon S3 bucket. To test the example, we recommend that you upload just a few documents to your S3 bucket. If your configuration is correct, you should eventually see a **Sync status** of **Completed**.

If you encounter any errors, see [Troubleshooting Amazon VPC connection](#).

Connecting to a database in a VPC

The following example shows how to connect a MySQL database running in a virtual private cloud (VPC). The example assumes that you're starting with your default VPC and that you need to create a MySQL database. If you already have a VPC, make sure that it's configured as shown. If you have a MySQL database, you can use that instead of creating a new one.

Steps

- [Step 1: Configure a VPC](#)
- [Step 2: Create and configure security groups](#)
- [Step 3: Create a database](#)
- [Step 4: Create a data source connector](#)

Step 1: Configure a VPC

Configure your VPC so that you have a private subnet and a security group for Amazon Q Business to access a MySQL database running in the subnet. The subnets provided in the VPC configuration must be in the US West (Oregon) Region, the US East (N. Virginia) Region, or the Europe (Ireland) Region.

To configure a VPC using Amazon VPC

1. Sign in to the AWS Management Console and open the Amazon VPC console at <https://console.aws.amazon.com/vpc/>.
2. From the navigation pane, choose **Route tables**, then choose **Create route table**.
3. For the **Name** field, enter **Private subnet route table**. From the **VPC** dropdown, select your VPC, and then choose **Create route table**. Choose **Close** to return to the list of route tables.

4. From the navigation pane, choose **NAT gateways**, then choose **Create NAT gateway**.
5. From the **Subnet** dropdown, choose the subnet that's the public subnet. Make a note of the subnet ID.
6. If you don't have an Elastic IP address, choose **Create New EIP**, choose **Create a NAT Gateway**, and then choose **Close**.
7. From the navigation pane, choose **Route tables**.
8. From the route table list, choose the **Private subnet route table** that you created in step 3. From **Actions**, choose **Edit routes**.
9. Choose **Add route**. For the destination, enter **0.0.0.0/0** to allow all outgoing traffic to the internet. For **Target**, choose **NAT Gateway**, and then choose the gateway that you created in step 4. Choose **Save changes**, and then choose **Close**.
10. From **Actions**, choose **Edit subnet associations**.
11. Choose the subnets that you want to be private. Don't choose the subnet with the NAT gateway that you noted previously. Choose **Save associations** when you're done.

Step 2: Create and configure security groups

Next, configure security groups for your database.

To create and configure security groups

1. Sign in to the AWS Management Console and open the Amazon VPC console at <https://console.aws.amazon.com/vpc/>.
2. From the description of your VPC, note the IPv4 CIDR.
3. From the navigation pane, choose **Security groups** and then choose **Create security group**.
4. For **Security group name**, enter **DataSourceInboundSecurityGroup**. Provide a description, then choose your VPC from the list. Choose **Create security group** and then choose **Close**.
5. Choose the **Inbound rules** tab.
6. Choose **Edit inbound rules**, and then choose **Add rule**
7. For a database, enter the port number for the **Port range**. For example, for MySQL it's **3306**, and, for HTTPS, it's **443**. For the **Source**, type the Classless Inter-Domain Routing (CIDR) of your VPC. Choose **Save rules** and then choose **Close**.

The security group allows anyone within the VPC to connect to the database, and it allows outbound connections to the internet.

Step 3: Create a database

Create a database to hold your documents, or you can use your existing database.

For instructions on how to create a MySQL database, see [MySQL](#).

Step 4: Create a data source connector

After you configure your VPC and create your database, you can create a data source connector for the database. For information about database connectors that Amazon Q supports, see [Supported connectors](#).

For your database, make sure that you configure your VPC, the private subnets that you created in your VPC, and the security group that you created in your VPC.

For instructions on how to create a data source for a MySQL database, see [MySQL](#).

Troubleshooting VPC connection issues

If you encounter any issues with your virtual private cloud (VPC) connection, check that your IAM permissions, security group settings, and the subnet's route tables are configured correctly.

One potential cause of a failed data source connector sync is that the data source might be unreachable from the subnet that you assigned to Amazon Q Business. To troubleshoot this issue, we recommend that you create an Amazon EC2 instance with the same Amazon VPC settings. Then, try to access the data source from this Amazon EC2 instance using REST API calls or other methods (based on the specific type of your data source).

If you successfully access the data source from the Amazon EC2 instance that you create, it means your data source is reachable from this subnet. Therefore, your sync issue isn't related to your data source being inaccessible by Amazon VPC.

If you can't access your Amazon EC2 instance from your VPC configuration and validate it with the Amazon EC2 instance that you created, you need to troubleshoot further. For example, if you have an Amazon S3 connector whose sync failed with errors about connection issues, you can set up an Amazon EC2 instance with the same Amazon VPC configuration that you assigned to your

Amazon S3 connector. Then, use this Amazon EC2 instance to test if your Amazon VPC has been set up correctly.

The following is an example of setting up an Amazon EC2 instance to troubleshoot your Amazon VPC connection with an Amazon S3 data source.

Topics

- [Step 1: Launch an Amazon EC2 instance](#)
- [Step 2: Connect to Amazon EC2 instance](#)
- [Step 3: Test Amazon S3 access](#)

Step 1: Launch an Amazon EC2 instance

1. Sign in to the AWS Management Console and open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Select **Launch an instance**.
3. Choose **Network settings**, and then choose **Edit**, and then do the following:
 - a. Choose the same VPC and **Subnet** that you assigned to Amazon Q.
 - b. For **Firewall (security groups)**, choose **Select existing security group**. Then, select the security group that you assigned to Amazon Q.

Note

The security group should allow outbound traffic to Amazon S3.

- c. Set **Auto-assign public IP** to **Disable**.
- d. In **Advanced details**, do the following:
 - For **IAM instance profile**, select **Create new IAM profile** to create and attach an IAM instance profile to your instance. Make sure that the profile has permissions to access Amazon S3. For more information, see [How can I grant my Amazon EC2 instance access to an Amazon S3 bucket?](#) in AWS re:Post.
 - Leave all other settings as default.
- e. Review and launch the Amazon EC2 instance.

Step 2: Connect to Amazon EC2 instance

After your Amazon EC2 instance is running, go to your instance detail page and connect to your instance. To do so, use the steps in [Connect to your instances without requiring a public IPv4 address using EC2 Instance Connect Endpoint](#) in the *Amazon EC2 User Guide for Linux Instances*.

Step 3: Test Amazon S3 access

After you have connected to your Amazon EC2 instance terminal, run an AWS CLI command to test the connection from this private subnet to your Amazon S3 bucket.

To test Amazon S3 access, type the following AWS CLI command in the AWS CLI: `aws s3 ls`

After the AWS CLI command runs, review the following:

- If you've set up the necessary IAM permissions correctly and your Amazon S3 setup is correct, you should see a list of your Amazon S3 buckets.
- If you see permission errors such as `Access Denied`, it's likely that your VPC configuration is correct, but something is wrong with your IAM permissions or Amazon S3 bucket policy.

If the command is timing out, then it's likely that your connection is timing out because your VPC setup is incorrect and the Amazon EC2 instance can't access Amazon S3 from your subnet. Reconfigure your VPC, and try again.

Troubleshooting data source connectors

This section can help you fix issues with Amazon Q Business data source connectors.

Topics

- [My documents were not indexed](#)
- [My synchronization job failed](#)
- [My synchronization job is incomplete](#)
- [My synchronization job succeeded but there are no indexed documents](#)
- [I am running into file format issues while syncing my data source](#)
- [I am getting an AccessDenied When Using SSL Certificate File error message](#)

My documents were not indexed

When you synchronize your Amazon Q Business index with a data source, you may run into issues that prevent the documents from being indexed. Indexing is a two-step process. First, the data source is checked for new and updated documents to index, and to find documents to remove from the index. Second, at the document level, each document is accessed and indexed.

An error can occur in either of these steps. Data source level errors are reported in the console in the **Sync run history** section of the data source details page. The status of the synchronization job can be **Succeeded**, **Incomplete**, or **Failed**. You can also see the number of documents indexed and deleted during the job. If the status is **Failed**, a message is shown in the **Details** column.

Document level errors are reported in Amazon CloudWatch Logs. You can see the errors using the CloudWatch console.

My synchronization job failed

A synchronization job typically fails when there is a configuration error in the index or the data source. In the console, you can find the error message in the **Sync run history** section of the data source details page, under the **Details** column. Document level errors are reported in Amazon CloudWatch Logs. The error message gives information about what went wrong. The problem is usually that the index or the data source doesn't have the correct IAM permissions. The error message describes the missing permissions. Following are some of the error messages that you can receive:

```
Failed to create log group for job. Please make sure that the IAM role provided has sufficient permissions.
```

If your index role doesn't have permissions to use CloudWatch, the data source can't create a CloudWatch log. If you get this error, you must add CloudWatch permissions to the index role.

```
Failed to access Amazon S3 file prefix (bucket name) while trying to crawl your metadata files. Please make sure the IAM role (ARN) provided has sufficient permissions.
```

When you're using an Amazon S3 data source, Amazon Q Business must have permissions to access the bucket that contains the documents. You need to add permissions for Amazon Q to read the bucket to the data source IAM role.

The provided IAM role (*ARN*) could not be assumed. Please make sure Amazon Q is a trusted entity that is allowed to assume the role.

Amazon Q needs permissions to assume the index and data source IAM roles. You need to add a trust policy to the roles with permissions for the `sts:AssumeRole` action.

For the IAM policies that Amazon Q needs to index a data source, see [IAM roles for Amazon Q Business connectors](#).

My synchronization job is incomplete

Jobs are generally incomplete when they have completed the data source level process but have some error during the document level process. When a job is incomplete, some of the documents might not have indexed successfully. For an Amazon S3 data source, an incomplete job is typically caused by one of the following issues:

- The metadata for one or more documents was not valid.
- When documents are submitted for indexing but at least one document was not submitted.
- When documents are submitted for deleting from the index but at least one document was not submitted.

To troubleshoot an incomplete synchronization job, look first to your CloudWatch logs.

1. From the details column, choose **View details in CloudWatch**.
2. Review the error messages to see what caused the document to fail.

My synchronization job succeeded but there are no indexed documents

Occasionally, an index synchronization job run is marked as **Succeeded**, but there are no new or updated documents indexed when you expect them. Possible reasons include the following:

- Check CloudWatch `DocumentsSubmittedForIndexingFailed` metric to see if any documents failed to synchronize. Check your CloudWatch logs for details.
- For an Amazon S3 data source, you might have given Amazon Q Business the wrong bucket name or prefix. Make sure that the S3 bucket that Amazon Q is using is the bucket that contains the documents to index.

- When re-indexing a document that failed to be indexed in an earlier job, Amazon Q won't index it unless you've changed the document or its associated metadata file.

I am running into file format issues while syncing my data source

If you run into file format issues while adding files to your data source or syncing your data source, make sure that your document types are supported by Amazon Q Business. For a list of document types supported by Amazon Q see [Supported document types](#).

If you're using the BatchPutDocument API operation with plaintext files, specify PLAIN_TEXT as the content type.

I am getting an AccessDenied When Using SSL Certificate File error message

If you're getting an access denied error when using an SSL certificate with your data source, make sure that your IAM role has the permissions to access the SSL certificate file in its specified location. If the certificate is encrypted with an AWS KMS key, your IAM role should also have permissions to decrypt using the AWS KMS key. For more information, see [Authentication and access control for AWS KMS](#) in the *AWS Key Management Service Developer Guide*.

Enhancing an Amazon Q Business application

After you finish [configuring your application](#), you can optionally choose to enhance it.

You can choose from the following available enhancements:

- **Document enrichment** – Control document attribute ingestion and build customized data solutions.
- **Guardrails** – Customize blocked topics and choose the knowledge sources your web experience uses for responses.
- **Plugins** – Enable your end users to perform specific tasks related to third-party services from within their web experience chat—like creating Jira tickets.
- **Relevance tuning** – Use document attributes to boost response generation from specific content within your application.

Topics

- [Document enrichment in Amazon Q Business](#)
- [Configuring plugins with Amazon Q Business](#)
- [Admin controls and guardrails in Amazon Q Business](#)
- [Boosting chat responses using relevance tuning](#)

Document enrichment in Amazon Q Business

The Amazon Q Business *document enrichment* feature helps you control both **what** documents and document attributes are ingested into your index and also **how** they're ingested. Using document enrichment, you can create, modify, or delete document attributes and document content when you ingest them into your Amazon Q index.

Document enrichment offers two kinds of methods that you can use for your solution:

- **Configure basic operations** – Use basic operations to add, update, or delete document attributes from your data. For example, you can scrub personally identifiable information (PII) by choosing to delete any document attributes related to PII.
- **Configure Lambda functions** – Use a preconfigured Lambda function to perform more customized, advanced document attribute manipulation logic to your data. For example, your

enterprise data might be stored as scanned images. In that case, you can use a Lambda function to run Optical Character recognition (OCR) on the scanned documents to extract text from them. Then, each scanned document is treated as a text document during ingestion. Finally, during chat, Amazon Q will factor the textual data extracted from the scanned documents when it generates responses.

When you implement your solution, you can choose to use both document enrichment methods together. That is, you can use basic operations to do a first parse of your data and then use a Lambda function for more complex operations. For example, you could first use a basic function to remove all PII information from your documents using document attributes. Then, use a Lambda function to extract text from scanned documents.

Document enrichment is supported both on the AWS Management Console and by Amazon Q API actions. If you use the console, you can only enrich documents connected to your application using an Amazon Q data source.

Note

Document enrichment is only supported in an Amazon Q application if you use an Amazon Q native retriever. If you use an Amazon Kendra retriever, you should [configure document enrichment](#) in Amazon Kendra.

Topics

- [How document enrichment works](#)
- [Using basic operations for document enrichment](#)
- [Using Lambda functions](#)

How document enrichment works

To understand and use document enrichments, you should be familiar with the key Amazon Q Business concepts that this topic outlines.

Topics

- [Document enrichment concepts](#)
- [Document enrichment process overview](#)

Document enrichment concepts

Amazon Q extracts *document attributes* from any document that you ingest into an Amazon Q index. Document attributes or structural metadata can include document title, document type, and time and date created. You can map document attributes to fields in an Amazon Q index to better structure your data for retrieval and chat. For more information, see [Document attributes and types](#) and [Filtering using document attributes](#).

Note

Although document attributes and index fields are distinct concepts, in practice they're used interchangeably because their values overlap and they structurally correspond to each other. That is, document attributes == document metadata == index fields.

Document enrichment process overview

The overall process of document enrichment is as follows:

- You configure document enrichment when you create or update your Amazon Q data source, or add or upload your documents directly into Amazon Q index. The exact process for configuration depends on the methods you choose:
 - If you use the API and want to configure document enrichment for a data source connector, you use the [CreateDataSource](#) and [UpdateDataSource](#) operations to provide your configuration details.
 - If you use the API and choose to directly upload documents into your index using the [BatchPutDocument](#) operation, you must configure document enrichment with each request.
 - If you use the console, can only configure document enrichment for a data source connected to your Amazon Q application. You select **Document enrichments** under **Enhancements** from the left navigation pane and configure enrichments. You can choose to use both configuration options or either one. You can also choose whether you want to apply your configuration to the original pre-extraction data or to the structured post-extraction data.
- After you configure and activate your document enrichment configuration, you can use inline configuration or basic logic to alter your data. For more information, see [Using basic operations](#).
- If you chose to configure advanced data manipulation by using a Lambda function, Amazon Q applies the configured function (depending on what you've chosen) to either your original pre-

extraction data or your structured post-extraction data. For more information, see [Using Lambda functions](#).

- Finally, your altered and enriched documents are ingested into your Amazon Q index.

If a configuration isn't valid during any point in this process, Amazon Q returns an error.

Using basic operations for document enrichment

With document enrichment, you can use basic operations to manipulate document attributes. For example, you can remove document attribute values, modify attribute values using conditions, or create document attributes.

Note

Amazon Q Business can't create a target document attribute field if it isn't already created as an index field.

Topics

- [Basic operations using the Amazon Q Business API](#)
- [Basic operations using the Amazon Q Business console](#)
- [Use cases for basic operations](#)
- [Code examples of basic operations](#)

Basic operations using the Amazon Q Business API

To apply basic logic, you specify your document attribute configuration using the [DocumentAttributeTarget](#) object when you use either the [BatchPutDocument](#) API operation or the [CreateDataSource](#) operation. Use the following parameters to create your configuration:

- `key` – The target field that you want to manipulate. For example, the key `Department` is a field or attribute that holds all the department names associated with the documents.
- `value` – The target value for your target attribute.
- `attributeValueOperator` – To delete an existing target value, set to `DELETE`. The default value for this parameter is `UPDATE`.

If a specific condition is met, you can also specify a value to use in the target field. Set the condition using the [DocumentAttributeCondition](#) object. For example, if the `_source_uri` field contains `financial` in its URI value, you can choose to prefill the target field `department` with the target value `finance` for the document.

For more information, see the following topics in the *Amazon Q API Reference*:

- [BatchPutDocument](#)
- [CreateDataSource](#)
- [DocumentAttributeTarget](#)
- [DocumentAttributeCondition](#)

Basic operations using the Amazon Q Business console

To apply basic logic using the console

1. Sign in to the AWS Management Console and open the Amazon Q console.
2. In **Applications**, select the name of your application from the list of applications.
3. From the left navigation menu, choose **Enhancements**, and then choose **Document enrichments**.
4. In **Document enrichments**, choose **Add document enrichment**.
5. In **Configure basic operations**, for **Document enrichment source**, choose a data source connected to your application.
6. To apply basic manipulations to your document fields and content, go to **Configure basic operations**.
7. Choose **Next** to save your configuration.

Use cases for basic operations

This section provides two examples of basic operations.

Example 1: Removing customer identification numbers associated with the documents

The following is an example of using a basic operation to remove all customer identification numbers in the document field called `customer_id`.

The following table shows the data before basic manipulation is applied.

_document_id	_document_id	customer_id
1	Example text	CID1234
2	Example text	CID1235
3	Example text	CID1236

The following table shows the data after basic manipulation is applied.

_document_id	_document_body	customer_id
1	Example text	
2	Example text	
3	Example text	

Example 2: Creating and prefilling the Department field with department names associated with the documents using a condition

The following is an example of using basic logic to create a field called Department and prefilling the field with the department names based on information from the `_source_uri` field. This example uses the condition that, if the `_source_uri` field contains `financial` in its URI value, then the target field `department` is prefilled with the target value `finance` for the document.

The following table shows the data before basic manipulation is applied.

_document_id	document_body	_source_uri
1	Example text	financial/1
2	Example text	financial/2
3	Example text	financial/3

The following table shows the data after basic manipulation is applied.

<code>_document_id</code>	<code>_document_body</code>	<code>_source_uri</code>	<code>department</code>
1	Example text	financial/1	Finance
2	Example text	financial/2	Finance
3	Example text	financial/3	Finance

Code examples of basic operations

The following instructions give examples of configuring basic data manipulation to remove customer identification numbers associated with the documents.

Console

To configure basic data manipulation to remove customer identification numbers

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. From the left navigation pane, select **Document enrichments** and then select **Add document enrichment**.
3. On the **Configure basic operations** page, choose from the data source that you want to alter document fields and content in.
4. Select the document field name **Customer_ID** from the dropdown menu, and then select the target action **Delete**.
5. Select **Add basic operation**.

AWS CLI

To configure basic data manipulation to remove customer identification numbers

```
aws qbusiness create-data-source \  
  --name data-source-name \  
  --application-id application-id \  
  --index-id index-id \  
  --role-arn arn:aws:iam::account-id:role/role-name \  
  --type S3 \  
  --configuration '{"S3Configuration":{"BucketName":"S3-bucket-name"}}' \  
  \
```

```
--document-enrichment-configuration '{"InlineDocumentEnrichmentConfiguration": [{"Target":{"key":"Customer_ID", "attributeValueOperator": "DELETE"}}]}'
```

Python

To configure basic data manipulation to remove customer identification numbers

```
import boto3
from botocore.exceptions import ClientError
import pprint
import time

qbusiness = boto3.client("qbusiness")

print("Create a data source with customizations")

# Provide the name of the data source
name = "data-source-name"
# Provide the application ID for the data source
application_id = "application-id"
# Provide the index ID for the data source
index_id = "index-id"
# Provide the IAM role ARN required for data sources
role_arn = "arn:aws:iam:${account-id}:role/${role-name}"
# Provide the data source connection information
data_source_type = "S3"
S3_bucket_name = "S3-bucket-name"
# Configure the data source with Document Enrichment
configuration = {"S3Configuration":
    {
        "BucketName": S3_bucket_name
    }
}
document_enrichment_configuration = {"InlineDocumentEnrichmentConfiguration": [
    {
        "Target":{"key":"Customer_ID",
            "attributeValueOperator": "DELETE"}
    }
]}

try:
    data_source_response = qbusiness.create_data_source(
        Name = name,
        ApplicationId = application_id,
```



```
        IndexId = index_id,
        RoleArn = role_arn,
        Type = data_source_type
        Configuration = configuration
        DocumentEnrichmentConfiguration = document_enrichment_configuration
    )

pprint.pprint(data_source_response)

data_source_id = data_source_response["Id"]

print("Wait for Amazon Q to create the data source with your customizations.")

while True:
    # Get the details of the data source, such as the status
    data_source_description = qbusiness.get_data_source(
        DataSourceId = data_source_id,
        ApplicationId = application_id,
        IndexId = index_id
    )
    status = data_source_description["Status"]
    print(" Creating data source. Status: "+status)
    time.sleep(60)
    if status != "CREATING":
        break

print("Synchronize the data source.")

sync_response = qbusiness.start_data_source_sync_job(
    DataSourceId = data_source_id,
    ApplicationId = application_id,
    IndexId = index_id
)

pprint.pprint(sync_response)

print("Wait for the data source to sync with the index.")

while True:

    jobs = qbusiness.list_data_source_sync_jobs(
        DataSourceId = data_source_id,
        ApplicationId = application_id,
        IndexId = index_id
```

```
)

# For this example, there should be one job
status = jobs["History"][0]["Status"]

print(" Syncing data source. Status: "+status)
time.sleep(60)
if status != "SYNCING":
    break

except ClientError as e:
    print("%s" % e)

print("Program ends.")
```

Java

To configure basic data manipulation to remove customer identification numbers

```
package com.amazonaws.qbusiness;

import java.util.concurrent.TimeUnit;
import software.amazon.awssdk.services.qbusiness.QBusinessClient;
import software.amazon.awssdk.services.qbusiness.model.AttributeValueOperator;
import software.amazon.awssdk.services.qbusiness.model.CreateDataSourceRequest;
import software.amazon.awssdk.services.qbusiness.model.CreateDataSourceResponse;
import software.amazon.awssdk.services.qbusiness.model.CreateIndexRequest;
import software.amazon.awssdk.services.qbusiness.model.CreateIndexResponse;
import software.amazon.awssdk.services.qbusiness.model.DataSourceConfiguration;
import software.amazon.awssdk.services.qbusiness.model.DataSourceStatus;
import software.amazon.awssdk.services.qbusiness.model.DataSourceSyncJob;
import software.amazon.awssdk.services.qbusiness.model.DataSourceSyncJobStatus;
import software.amazon.awssdk.services.qbusiness.model.DataSourceType;
import software.amazon.awssdk.services.qbusiness.model.GetDataSourceRequest;
import software.amazon.awssdk.services.qbusiness.model.GetDataSourceResponse;
import software.amazon.awssdk.services.qbusiness.model.IndexStatus;
import
    software.amazon.awssdk.services.qbusiness.model.ListDataSourceSyncJobsRequest;
import
    software.amazon.awssdk.services.qbusiness.model.ListDataSourceSyncJobsResponse;
import software.amazon.awssdk.services.qbusiness.model.DataSourceConfiguration;
import
    software.amazon.awssdk.services.qbusiness.model.StartDataSourceSyncJobRequest;
```

```
import
software.amazon.awssdk.services.qbusiness.model.StartDataSourceSyncJobResponse;

public class CreateDataSourceWithCustomizationsExample {

    public static void main(String[] args) throws InterruptedException {
        System.out.println("Create a data source with customizations");

        String dataSourceName = "data-source-name";
        String applicationId = "application-id";
        String indexId = "index-id";
        String dataSourceRoleArn = "arn:aws:iam::account-id:role/role-name";
        String s3BucketName = "S3-bucket-name"

        QBusinessClient qbusiness = QBusinessClient.builder().build();

        CreateDataSourceRequest createDataSourceRequest = CreateDataSourceRequest
            .builder()
            .name(dataSourceName)
            .applicationId(applicationId)
            .indexId(indexId)
            .description(experienceDescription)
            .roleArn(experienceRoleArn)
            .type(DataSourceType.S3)
            .configuration(
                DataSourceConfiguration
                    .builder()
                    .s3Configuration(
                        S3DataSourceConfiguration
                            .builder()
                            .bucketName(s3BucketName)
                            .build()
                    ).build()
            )
            .documentEnrichmentConfiguration(
                DocumentEnrichmentConfiguration
                    .builder()
                    .inlineDocumentEnrichmentConfiguration(Arrays.asList(
                        InlineDocumentEnrichmentConfiguration
                            .builder()
                            .target(
                                DocumentAttributeTarget
                                    .builder()
                                    .key("Customer_ID")
```

```
.attributeValueOperator(AttributeValueOperator.DELETE)
    .build()
    .build()
    ).build();

CreateDataSourceResponse createDataSourceResponse =
qbusiness.createDataSource(createDataSourceRequest);
    System.out.println(String.format("Response of creating data source: %s",
createDataSourceResponse));

String dataSourceId = createDataSourceResponse.id();
System.out.println(String.format("Waiting for Amazon Q to create the data
source %s", dataSourceId));
GetDataSourceRequest getDataSourceRequest = GetDataSourceRequest
    .builder()
    .applicationId(applicationId)
    .indexId(indexId)
    .datasourceId(dataSourceId)
    .build();

while (true) {
    GetDataSourceResponse getDataSourceResponse =
qbusiness.getDataSource(getDataSourceRequest);

    DataSourceStatus status = getDataSourceResponse.status();
    System.out.println(String.format("Creating data source. Status: %s",
status));
    TimeUnit.SECONDS.sleep(60);
    if (status != DataSourceStatus.CREATING) {
        break;
    }
}

System.out.println(String.format("Synchronize the data source %s",
dataSourceId));
StartDataSourceSyncJobRequest startDataSourceSyncJobRequest =
StartDataSourceSyncJobRequest
    .builder()
    .applicationId(applicationId)
    .indexId(indexId)
    .datasourceId(dataSourceId)
    .build();
```

```
        StartDataSourceSyncJobResponse startDataSourceSyncJobResponse =
qbusiness.startDataSourceSyncJob(startDataSourceSyncJobRequest);
        System.out.println(String.format("Waiting for the data source to sync
with the application %s index %s for execution ID %s", applicationId, indexId,
startDataSourceSyncJobResponse.executionId()));

        // For this example, there should be one job
        ListDataSourceSyncJobsRequest listDataSourceSyncJobsRequest =
ListDataSourceSyncJobsRequest
            .builder()
            .applicationId(applicationId)
            .indexId(indexId)
            .datasourceId(dataSourceId)
            .build();

        while (true) {
            ListDataSourceSyncJobsResponse listDataSourceSyncJobsResponse =
qbusiness.listDataSourceSyncJobs(listDataSourceSyncJobsRequest);
            DataSourceSyncJob job = listDataSourceSyncJobsResponse.history().get(0);
            System.out.println(String.format("Syncing data source. Status: %s",
job.status()));

            TimeUnit.SECONDS.sleep(60);
            if (job.status() != DataSourceSyncJobStatus.SYNCING) {
                break;
            }
        }

        System.out.println("Data source creation with customizations is complete");
    }
}
```

Using Lambda functions

You can use Lambda functions to prepare your document attributes for advanced data manipulation. For example, you could use Optical Character Recognition (OCR), which interprets text from images and treats each image as a textual document. Or, you could retrieve the current date-time in a specific time zone and then insert the date-time where there's an empty value for a date field.

You can choose to apply a basic operation first and then use a Lambda function to manipulate your data, and the reverse.

Note

Amazon Q Business can't create a target document attribute field if it isn't already created as an index field.

Topics

- [Lambda functions using the Amazon Q Business API](#)
- [Lambda functions using the Amazon Q Business console](#)
- [IAM roles for Lambda functions](#)
- [Use cases for Lambda functions](#)
- [Code examples of Lambda functions](#)
- [Data contracts for Lambda functions](#)

Lambda functions using the Amazon Q Business API

To apply a Lambda function, you specify your advanced data manipulation logic using the [DocumentEnrichmentConfiguration](#) object when you use either the [BatchPutDocument](#) API operation or the [CreateDataSource](#) operation.

Your Lambda functions must follow the mandatory request and response structures. For more information, see [Data contracts for Lambda functions](#).

Use the following parameters to create your configuration:

- `InlineDocumentEnrichmentConfiguration` – Configuration information to alter document attributes during ingestion.
- `PostExtractionHookConfiguration` – Configuration information to invoke a Lambda function on structured documents with their metadata and text already extracted.
- `PreExtractionHookConfiguration` – Configuration information to invoke a Lambda function on raw documents before metadata and text has been extracted from them.
- `PreExtractionHookConfiguration RoleArn` – The Amazon Resource Name (ARN) of a role under `PreExtractionHookConfiguration` with permissions to run

`PreExtractionHookConfiguration` and to access the Amazon S3 bucket when you use `PreExtractionHookConfiguration`.

- `PostExtractionHookConfiguration RoleArn` – The Amazon Resource Name (ARN) of a role under `PostExtractionHookConfiguration` with permissions to run `PreExtractionHookConfiguration` and to access the Amazon S3 bucket when you use `PostExtractionHookConfiguration`.

You can configure only one Lambda function for `PreExtractionHookConfiguration` and only one Lambda function for `PostExtractionHookConfiguration`. However, your Lambda function can invoke other functions that it requires.

You can configure both `PreExtractionHookConfiguration` and `PostExtractionHookConfiguration` or either one. Your Lambda function for `PreExtractionHookConfiguration` must not exceed a run time of 5 minutes. Your Lambda function for `PostExtractionHookConfiguration` must not exceed a run time of 1 minute.

You can configure Amazon Q to invoke a Lambda function only if a condition is met. For example, you can specify a condition that, if there are empty date-time values, then Amazon Q invokes a function that inserts the current date-time.

For more information, see the following topics in the *Amazon Q API Reference*:

- [BatchPutDocument](#)
- [CreateDataSource](#)
- [DocumentEnrichmentConfiguration](#)
- [DocumentAttributeCondition](#)

Lambda functions using the Amazon Q Business console

To configure a Lambda function using the console

1. Select your index, and then select **Document enrichments** from the navigation menu.
2. To configure Lambda functions, go to **Configure Lambda functions**.

IAM roles for Lambda functions

When you use the Lambda functions for CDE, you need an IAM role for the following:

- A role for `PreExtractionHookConfiguration` with permissions to run `PreExtractionHookConfiguration` and to access the Amazon S3 bucket when you use `PreExtractionHookConfiguration`.
- A role for `PostExtractionHookConfiguration` with permissions to run `PreExtractionHookConfiguration` and to access the Amazon S3 bucket when you use `PostExtractionHookConfiguration`.

Important

IAM roles for Custom Document Enrichment (CDE) Lambda functions should belong to the same account as the account using [BatchPutDocument](#) API operation or the [CreateDataSource](#) operation to configure CDE.

Both AWS Identity and Access Management (IAM) roles must have the permissions to:

- Run `PreExtractionHookConfiguration` and/or `PostExtractionHookConfiguration`. To apply advanced alterations of your document metadata and content during the ingestion process, configure a Lambda function for `PreExtractionHookConfiguration` and/or `PostExtractionHookConfiguration`.
- (Optional) If you choose to activate Server Side Encryption for your Amazon S3 bucket, you must provide permissions to use the AWS KMS key to encrypt and decrypt the objects stored in your Amazon S3 bucket.

A role policy to allow Amazon Q to run `PreExtractionHookConfiguration` with encryption for your Amazon S3 bucket.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Action": [
      "s3:GetObject",
      "s3:PutObject"
    ],
    "Resource": [
      "arn:aws:s3:::bucket-name",
      "arn:aws:s3:::bucket-name/*"
    ],
  ]
}
```



```

    "Effect": "Allow"
  },
  {
    "Action": [
      "s3:ListBucket"
    ],
    "Resource": [
      "arn:aws:s3:::bucket-name"
    ],
    "Effect": "Allow"
  },
  {
    "Effect": "Allow",
    "Action": [
      "kms:Decrypt",
      "kms:GenerateDataKey"
    ],
    "Resource": [
      "arn:aws:kms:your-region:your-account-id:key/key-id"
    ]
  },
  {
    "Effect": "Allow",
    "Action": [
      "lambda:InvokeFunction"
    ],
    "Resource": "arn:aws:lambda:your-region:your-account-id:function:pre-
extraction-lambda-function"
  }
]
}

```

An role policy to allow Amazon Q to run PreExtractionHookConfiguration without encryption.

```

{
  "Version": "2012-10-17",
  "Statement": [{
    "Action": [
      "s3:GetObject",
      "s3:PutObject"
    ],
    "Resource": [

```

```

        "arn:aws:s3:::bucket-name",
        "arn:aws:s3:::bucket-name/*"
    ],
    "Effect": "Allow"
},
{
    "Action": [
        "s3:ListBucket"
    ],
    "Resource": [
        "arn:aws:s3:::bucket-name"
    ],
    "Effect": "Allow"
},
{
    "Effect": "Allow",
    "Action": [
        "lambda:InvokeFunction"
    ],
    "Resource": "arn:aws:lambda:your-region:your-account-id:function:pre-
extraction-lambda-function"
}
]
}

```

A role policy to allow Amazon Q to run `PostExtractionHookConfiguration` with encryption for your Amazon S3 bucket.

```

{
    "Version": "2012-10-17",
    "Statement": [{
        "Action": [
            "s3:GetObject",
            "s3:PutObject"
        ],
        "Resource": [
            "arn:aws:s3:::bucket-name",
            "arn:aws:s3:::bucket-name/*"
        ],
        "Effect": "Allow"
    },
    {
        "Action": [

```

```

        "s3:ListBucket"
    ],
    "Resource": [
        "arn:aws:s3:::bucket-name"
    ],
    "Effect": "Allow"
},
{
    "Effect": "Allow",
    "Action": [
        "kms:Decrypt",
        "kms:GenerateDataKey"
    ],
    "Resource": [
        "arn:aws:kms:your-region:your-account-id:key/key-id"
    ]
},
{
    "Effect": "Allow",
    "Action": [
        "lambda:InvokeFunction"
    ],
    "Resource": "arn:aws:lambda:your-region:your-account-id:function:post-
extraction-lambda-function"
}
]
}

```

An role policy to allow Amazon Q to run PostExtractionHookConfiguration without encryption.

```

{
  "Version": "2012-10-17",
  "Statement": [{
    "Action": [
      "s3:GetObject",
      "s3:PutObject"
    ],
    "Resource": [
      "arn:aws:s3:::bucket-name",
      "arn:aws:s3:::bucket-name/*"
    ],
  ]
}

```

```

    "Effect": "Allow"
  },
  {
    "Action": [
      "s3:ListBucket"
    ],
    "Resource": [
      "arn:aws:s3:::bucket-name"
    ],
    "Effect": "Allow"
  },
  {
    "Effect": "Allow",
    "Action": [
      "lambda:InvokeFunction"
    ],
    "Resource": "arn:aws:lambda:your-region:your-account-id:function:post-extraction-
lambda-function"
  ]
}

```

We recommend that you include `aws:sourceAccount` and `aws:sourceArn` in the trust policy. Their inclusion limits permissions and securely checks if `aws:sourceAccount` and `aws:sourceArn` are the same values as provided in the IAM role policy for the `sts:AssumeRole` action. This approach prevents unauthorized entities from accessing your IAM roles and their permissions. For more information, see [confused deputy problem](#) in the *IAM User Guide*.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": [
          "qbusiness.amazonaws.com"
        ]
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "your-account-id"
        },
        "StringLike": {

```

```

        "aws:SourceArn": "arn:aws:qbusiness:your-region:your-account-id:application/
<application-id>/index/<index-id>"
    }
}
]
}

```

Use cases for Lambda functions

This section outlines two examples of using Lambda functions.

Example 1: Extracting text from images to create textual documents

The following is an example of using a Lambda function to run OCR to interpret text from images and store this text in a field called `document_image_text`.

The following table shows data before advanced manipulation is applied.

<code>_document_id</code>	<code>document_image</code>
1	image_1.png
2	image_2.png
3	image_3.png

The following table shows data after advanced manipulation is applied.

<code>_document_id</code>	<code>document_image</code>	<code>document_image_text</code>
1	image_1.png	Mailed survey response
2	image_2.png	Mailed survey response
3	image_3.png	Mailed survey response

Example 2: Replacing empty values in the `Last_Updated` field with the current date-time

The following is an example of using a Lambda function to insert the current date-time for empty date values. This example uses the condition that, if a date field value is `null`, then the value is replaced with the current date-time.

The following table shows data before advanced manipulation is applied.

<code>_document_id</code>	<code>_document_body</code>	<code>_last_updated_at</code>
1	Example text	January 1, 2020
2	Example text	
3	Example text	July 1, 2020

The following table shows data after advanced manipulation is applied.

<code>_document_id</code>	<code>_document_body</code>	<code>_last_updated_at</code>
1	Example text	January 1, 2020
2	Example text	December 1, 2021
3	Example text	July 1, 2020

Code examples of Lambda functions

The following code is an example of configuring a Lambda function for advanced data manipulation on the raw, original data.

Console

To configure a Lambda function for advanced data manipulation on the raw, original data

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. From the left navigation menu, choose **Enhancements**, and then choose **Document enrichments**.
3. In **Document enrichments**, choose **Add document enrichment**.

4. In **Configure basic operations**, for **Document enrichment source**, choose a data source connected to your application.
5. (Optional) To apply basic manipulations to your document fields and content, go to **Configure basic operations** and choose **Next** to save your configuration.
6. On the **Configure Lambda functions** page, in the **Lambda for pre-extraction** section, select your Lambda function ARN and your Amazon S3 bucket using the dropdown menus.
7. To add your IAM access role, select the option to create a new role from the dropdown. This step creates the required Amazon Q permissions to create the document enrichment.
8. Select **Add basic operation**.

AWS CLI

To configure a Lambda function for advanced data manipulation on the raw, original data

```
aws qbusiness create-data-source \
  --name data-source-name \
  --application-id application-id \
  --index-id index-id \
  --role-arn arn:aws:iam::account-id:role/role-name \
  --type S3 \
  --configuration '{"S3Configuration":{"BucketName":"S3-bucket-name"}}' \
  --document-enrichment-configuration '{"InlineDocumentEnrichmentConfiguration":
[{"Target":{"key":"Customer_ID", "attributeValueOperator": true}}]}'
```

Python

To configure a Lambda function for advanced data manipulation on the raw, original data

```
import boto3
from botocore.exceptions import ClientError
import pprint
import time

qbusiness = boto3.client("qbusiness")

print("Create a data source with customizations")

# Provide the name of the data source
name = "data-source-name"
# Provide the application ID for the data source
```

```
application_id = "application-id"
# Provide the index ID for the data source
index_id = "index-id"
# Provide the IAM role ARN required for data sources
role_arn = "arn:aws:iam:${account-id}:role/${role-name}"
# Provide the data source connection information
data_source_type = "S3"
S3_bucket_name = "S3-bucket-name"
# Configure the data source with Document Enrichment
configuration = {"S3Configuration":
    {
        "BucketName": S3_bucket_name
    }
}
document_enrichment_configuration = {"InlineDocumentEnrichmentConfiguration":[
    {
        "Target":{"key":"Customer_ID",
            "attributeValueOperator": "DELETE"}
    }
]}

try:
    data_source_response = qbusiness.create_data_source(
        Name = name,
        ApplicationId = application_id,
        IndexId = index_id,
        RoleArn = role_arn,
        Type = data_source_type
        Configuration = configuration
        DocumentEnrichmentConfiguration = document_enrichment_configuration
    )

    pprint.pprint(data_source_response)

    data_source_id = data_source_response["Id"]

    print("Wait for Amazon Q to create the data source with your customizations.")

    while True:
        # Get the details of the data source, such as the status
        data_source_description = qbusiness.get_data_source(
            DataSourceId = data_source_id,
            ApplicationId = application_id,
            IndexId = index_id
```



```
)
    status = data_source_description["Status"]
    print(" Creating data source. Status: "+status)
    time.sleep(60)
    if status != "CREATING":
        break

print("Synchronize the data source.")

sync_response = qbusiness.start_data_source_sync_job(
    DataSourceId = data_source_id,
    ApplicationId = application_id,
    IndexId = index_id
)

pprint.pprint(sync_response)

print("Wait for the data source to sync with the index.")

while True:

    jobs = qbusiness.list_data_source_sync_jobs(
        DataSourceId = data_source_id,
        ApplicationId = application_id,
        IndexId = index_id
    )

    # For this example, there should be one job
    status = jobs["History"][0]["Status"]

    print(" Syncing data source. Status: "+status)
    time.sleep(60)
    if status != "SYNCING":
        break

except ClientError as e:
    print("%s" % e)

print("Program ends.")
```

Java

To configure a Lambda function for advanced data manipulation on the raw, original data

```
package com.amazonaws.qbusiness;

import java.util.concurrent.TimeUnit;
import software.amazon.awssdk.services.qbusiness.QBusinessClient;
import software.amazon.awssdk.services.qbusiness.model.AttributeValueOperator;
import software.amazon.awssdk.services.qbusiness.model.CreateDataSourceRequest;
import software.amazon.awssdk.services.qbusiness.model.CreateDataSourceResponse;
import software.amazon.awssdk.services.qbusiness.model.CreateIndexRequest;
import software.amazon.awssdk.services.qbusiness.model.CreateIndexResponse;
import software.amazon.awssdk.services.qbusiness.model.DataSourceConfiguration;
import software.amazon.awssdk.services.qbusiness.model.DataSourceStatus;
import software.amazon.awssdk.services.qbusiness.model.DataSourceSyncJob;
import software.amazon.awssdk.services.qbusiness.model.DataSourceSyncJobStatus;
import software.amazon.awssdk.services.qbusiness.model.DataSourceType;
import software.amazon.awssdk.services.qbusiness.model.GetDataSourceRequest;
import software.amazon.awssdk.services.qbusiness.model.GetDataSourceResponse;
import software.amazon.awssdk.services.qbusiness.model.IndexStatus;
import
    software.amazon.awssdk.services.qbusiness.model.ListDataSourceSyncJobsRequest;
import
    software.amazon.awssdk.services.qbusiness.model.ListDataSourceSyncJobsResponse;
import software.amazon.awssdk.services.qbusiness.model.DataSourceConfiguration;
import
    software.amazon.awssdk.services.qbusiness.model.StartDataSourceSyncJobRequest;
import
    software.amazon.awssdk.services.qbusiness.model.StartDataSourceSyncJobResponse;

public class CreateDataSourceWithCustomizationsExample {

    public static void main(String[] args) throws InterruptedException {
        System.out.println("Create a data source with customizations");

        String dataSourceName = "data-source-name";
        String applicationId = "application-id";
        String indexId = "index-id";
        String dataSourceRoleArn = "arn:aws:iam::account-id:role/role-name";
        String s3BucketName = "S3-bucket-name"

        QBusinessClient qbusiness = QBusinessClient.builder().build();

        CreateDataSourceRequest createDataSourceRequest = CreateDataSourceRequest
            .builder()
            .name(dataSourceName)
```

```

        .applicationId(applicationId)
        .indexId(indexId)
        .description(experienceDescription)
        .roleArn(experienceRoleArn)
        .type(DataSourceType.S3)
        .configuration(
            DataSourceConfiguration
                .builder()
                .s3Configuration(
                    S3DataSourceConfiguration
                        .builder()
                        .bucketName(s3BucketName)
                        .build()
                ).build()
        )
        .documentEnrichmentConfiguration(
            DocumentEnrichmentConfiguration
                .builder()
                .inlineConfigurations(Arrays.asList(
                    InlineDocumentEnrichmentConfiguration
                        .builder()
                        .target(
                            DocumentAttributeTarget
                                .builder()
                                .key("Customer_ID")
                        )
                ))
        ).attributeValueOperator(AttributeValueOperator.DELETE)
        .build()
    ).build();

    CreateDataSourceResponse createDataSourceResponse =
qbusiness.createDataSource(createDataSourceRequest);
    System.out.println(String.format("Response of creating data source: %s",
createDataSourceResponse));

    String dataSourceId = createDataSourceResponse.id();
    System.out.println(String.format("Waiting for Amazon Q to create the data
source %s", dataSourceId));
    GetDataSourceRequest getDataSourceRequest = GetDataSourceRequest
        .builder()
        .applicationId(applicationId)
        .indexId(indexId)
        .datasourceId(dataSourceId)

```

```
        .build();

    while (true) {
        GetDataSourceResponse getDataSourceResponse =
qbusiness.getDataSource(getDataSourceRequest);

        DataSourceStatus status = getDataSourceResponse.status();
        System.out.println(String.format("Creating data source. Status: %s",
status));
        TimeUnit.SECONDS.sleep(60);
        if (status != DataSourceStatus.CREATING) {
            break;
        }
    }

    System.out.println(String.format("Synchronize the data source %s",
dataSourceId));
    StartDataSourceSyncJobRequest startDataSourceSyncJobRequest =
StartDataSourceSyncJobRequest
        .builder()
        .applicationId(applicationId)
        .indexId(indexId)
        .datasourceId(dataSourceId)
        .build();
    StartDataSourceSyncJobResponse startDataSourceSyncJobResponse =
qbusiness.startDataSourceSyncJob(startDataSourceSyncJobRequest);
    System.out.println(String.format("Waiting for the data source to sync
with the application %s index %s for execution ID %s", applicationId, indexId,
startDataSourceSyncJobResponse.executionId()));

    // For this example, there should be one job
    ListDataSourceSyncJobsRequest listDataSourceSyncJobsRequest =
ListDataSourceSyncJobsRequest
        .builder()
        .applicationId(applicationId)
        .indexId(indexId)
        .datasourceId(dataSourceId)
        .build();

    while (true) {
        ListDataSourceSyncJobsResponse listDataSourceSyncJobsResponse =
qbusiness.listDataSourceSyncJobs(listDataSourceSyncJobsRequest);
        DataSourceSyncJob job = listDataSourceSyncJobsResponse.history().get(0);
```

```

        System.out.println(String.format("Syncing data source. Status: %s",
job.status()));

        TimeUnit.SECONDS.sleep(60);
        if (job.status() != DataSourceSyncJobStatus.SYNCING) {
            break;
        }
    }

    System.out.println("Data source creation with customizations is complete");
}
}

```

Data contracts for Lambda functions

Lambda functions for advanced data manipulation interact with Amazon Q data contracts. The contracts are the mandatory request and response structures of your Lambda functions. If your Lambda functions don't follow these structures, then Amazon Q produces an error. Your Lambda function for `PreExtractionHookConfiguration` should use the following request structure:

```

{
  "version": <str>,
  "dataBlobStringEncodedInBase64": <str>, //In the case of a data blob
  "s3Bucket": <str>, //In the case of an S3 bucket
  "s3ObjectKey": <str>, //In the case of an S3 bucket
  "metadata": <Metadata>
}

```

The metadata structure, which includes the `DocumentAttribute` structure, is as follows:

```

{
  "attributes": [<DocumentAttribute>]
}

DocumentAttribute
{
  "name": <str>,
  "value": <DocumentAttributeValue>
}

```

```
DocumentAttributeValue
{
  "stringValue": <str>,
  "integerValue": <int>,
  "longValue": <long>,
  "stringListValue": list<str>,
  "dateValue": <str>
}
```

Your Lambda function for `PreExtractionHookConfiguration` must adhere to the following response structure:

```
{
  "version": <str>,
  "dataBlobStringEncodedInBase64": <str>, //In the case of a data blob
  "s3ObjectKey": <str>, //In the case of an S3 bucket
  "metadataUpdates": [<DocumentAttribute>]
}
```

Your Lambda function for `PostExtractionHookConfiguration` should expect the following request structure:

```
{
  "version": <str>,
  "s3Bucket": <str>,
  "s3ObjectKey": <str>,
  "metadata": <Metadata>
}
```

Your Lambda function for `PostExtractionHookConfiguration` must adhere to the following response structure:

```
PostExtractionHookConfiguration Lambda Response
{
  "version": <str>,
  "s3ObjectKey": <str>,
  "metadataUpdates": [<DocumentAttribute>]
}
```

Amazon Q uploads your structured document to the specified Amazon S3 bucket. The structured document follows this format:

```
QBusiness document

{
  "textContent": <TextContent>
}

TextContent
{
  "documentBodyText": <str>
}
```

Examples of Lambda functions that adhere to data contracts

This section provides examples of how to structure your Lambda functions that adhere to Amazon Q data contracts.

Example 1: A Lambda function that applies advanced manipulation to raw documents

The following Python code is an example of a Lambda function that applies advanced manipulation of the metadata fields `_authors`, `_document_title`, and the body content on the raw or original documents.

The following code example shows the case of the body content residing in an Amazon S3 bucket

```
import json
import boto3

s3 = boto3.client("s3")

# Lambda function for advanced data manipulation
def lambda_handler(event, context):
    # Get the value of "S3Bucket" key name or item from the given event input
    s3_bucket = event.get("s3Bucket")
    # Get the value of "S3ObjectKey" key name or item from the given event input
    s3_object_key = event.get("s3ObjectKey")

    content_object_before_DE = s3.get_object(Bucket = s3_bucket, Key = s3_object_key)
    content_before_DE = content_object_before_DE["Body"].read().decode("utf-8");
    content_after_DE = "DEInvolved " + content_before_DE

    # Get the value of "metadata" key name or item from the given event input
    metadata = event.get("metadata")
```

```

# Get the document "attributes" from the metadata
document_attributes = metadata.get("attributes")

s3.put_object(Bucket = s3_bucket, Key = "dummy_updated_qbusiness_document",
Body=json.dumps(content_after_DE))
return {
    "version": "v0",
    "s3objectKey": "dummy_updated_qbusiness_document",
    "metadataUpdates": [
        {"name": "_document_title", "value":
{"stringValue": "title_from_pre_extraction_lambda"}},
        {"name": "_authors", "value": {"stringListValue": ["author1", "author2"]}}
    ]
}

```

Example 2: A Lambda function that applies advanced manipulation to structured or parsed documents

The following Python code is an example of a Lambda function that applies advanced manipulation of the metadata fields `_authors`, `_document_title`, and the body content on the structured or parsed documents.

```

import json
import boto3
import time

s3 = boto3.client("s3")

# Lambda function for advanced data manipulation
def lambda_handler(event, context):

    # Get the value of "S3Bucket" key name or item from the given event input
    s3_bucket = event.get("s3Bucket")
    # Get the value of "S3ObjectKey" key name or item from the given event input
    s3_key = event.get("s3ObjectKey")
    # Get the value of "metadata" key name or item from the given event input
    metadata = event.get("metadata")
    # Get the document "attributes" from the metadata
    document_attributes = metadata.get("attributes")

    qbusiness_document_object = s3.get_object(Bucket = s3_bucket, Key = s3_key)
    qbusiness_document_string =
qbusiness_document_object['Body'].read().decode('utf-8')

```



```

qbusiness_document = json.loads(qbusiness_document_string)
qbusiness_document["textContent"]["documentBodyText"] = "Changing document body to
a short sentence."

s3.put_object(Bucket = s3_bucket, Key = "dummy_updated_qbusiness_document",
Body=json.dumps(qbusiness_document))

return {
    "version" : "v0",
    "s3objectKey": "dummy_updated_qbusiness_document",
    "metadataUpdates": [
        {"name": "_document_title", "value":{"stringValue":
"title_from_post_extraction_lambda"}},
        {"name": "_authors", "value":{"stringListValue":["author1", "author2"]}}
    ]
}

```

Example 3: Body content residing in a data blob

```

import json
import boto3
import base64

# Lambda function for advanced data manipulation
def lambda_handler(event, context):

    # Get the value of "dataBlobStringEncodedInBase64" key name or item from the given
    event input
    data_blob_string_encoded_in_base64 = event.get("dataBlobStringEncodedInBase64")
    # Decode the data blob string in UTF-8
    data_blob_string =
base64.b64decode(data_blob_string_encoded_in_base64).decode("utf-8")
    # Get the value of "metadata" key name or item from the given event input
    metadata = event.get("metadata")
    # Get the document "attributes" from the metadata
    document_attributes = metadata.get("attributes")

    new_data_blob = "This should be the modified data in the document by pre processing
lambda ".encode("utf-8")
    return {
        "version": "v0",
        "dataBlobStringEncodedInBase64":
base64.b64encode(new_data_blob).decode("utf-8"),

```

```
    "metadataUpdates": [  
      {"name": "_document_title", "value":  
{"stringValue": "title_from_pre_extraction_lambda"}},  
      {"name": "_authors", "value": {"stringListValue": ["author1", "author2"]}}  
    ]  
  }
```

Configuring plugins with Amazon Q Business

You can create and configure plugins for your Amazon Q Business application. Once configured, plugins can support actions—or write operations and instructions—that can help you boost end user productivity. End users can perform specific tasks related to third-party services from within their web experience chat such as creating a Jira ticket.

For example, your end user might be an IT representative whose Amazon Q chat requires the follow-up action of opening an incident in ServiceNow. They can request that Amazon Q create an incident in ServiceNow on their behalf without leaving their chat.

Amazon Q supports the following plugins and actions:

- **Jira** – Creating an issue
- **Salesforce** – Creating a case
- **ServiceNow** – Creating an incident
- **Zendesk** – Creating a ticket

Each Amazon Q plugin supports a single action. Each Amazon Q application can have up to 3 enabled plugins. No two plugins can be of the same type. Once activated, you can choose to deactivate, reactivate, edit, and delete plugins at any time. You can't customize plugins.

Topics

- [Configuring a Jira plugin](#)
- [Configuring a Salesforce plugin](#)
- [Configuring a ServiceNow plugin](#)
- [Configuring a Zendesk plugin](#)
- [Using Amazon Q Business plugins](#)
- [Managing Amazon Q Business plugins](#)

Configuring a Jira plugin

Jira is a project management tool that creates issues (tickets) for software development, product management, and bug tracking. If you're a Jira user, you can create an Amazon Q Business plugin to allow your end users to create Jira issues from within their web experience chat.

To create a Jira plugin, you need configuration information from your Jira instance to set up a connection between Amazon Q and Jira and allow Amazon Q to perform actions in Jira.

For more information on how to use plugins during your web experience chat, see [Using plugins](#).

Topics

- [Prerequisites](#)
- [Service access roles](#)
- [Creating a plugin](#)

Prerequisites

Before you configure your Amazon Q Jira plugin, you must do the following:

- Set up a new user in your Jira instance with scoped permissions for performing actions in Amazon Q.
- (Optional) [Create an API token](#) for the new user that you created.
- Note this user's Jira username and Jira account password (and optionally, their API token). You will need this basic authentication information for creating an AWS Secrets Manager secret during the plugin configuration process.
- Note the base URL of your Jira Cloud instance hosted by Atlassian. For example: `https://yourcompany.atlassian.net`.

Service access roles

To successfully connect Amazon Q to Jira, you need to give Amazon Q the following permission to access your Secrets Manager secret to get your Jira credentials. Amazon Q assumes this role to access your Jira credentials.

The following is the service access IAM role required:

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Action": [
      "secretsmanager:GetSecretValue"
    ],
    "Resource": [
      "arn:aws:secretsmanager:{{your-region}}:{{your-account-id}}:secret:
[[secret-id]]"
    ]
  }
]
```

If you use the console and choose to create a new IAM role, Amazon Q creates the role for you. If you use the console and choose to use an existing secret, or you use the API, make sure your IAM role contains these permissions.

Creating a plugin

To create a Jira plugin for your web experience chat, you can use the AWS Management Console or the [CreatePlugin](#) API operation. The following tabs provide a procedure to create a Jira plugin using the console and code examples for the AWS CLI.

Console

To create a Jira plugin

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, select the name of your application from the list of applications.
3. From the left navigation menu, choose **Enhancements**, and then choose **Plugins**.
4. For **Plugins**, choose **Add plugin**.
5. For **Add plugins**, choose **Jira**.
6. For **Jira**, enter the following information:
 - a. **Name, Plugin name** – A name for your Amazon Q plugin. The name can include hyphens (-), but not spaces, and can have a maximum of 1,000 alphanumeric characters.

- b. **Service access** – Choose **Create and add a new service role** or **Use an existing service role**. Make sure that your service role has the necessary permissions.
- c. **URL** – The base URL of your Jira Cloud instance hosted by Atlassian. For example: `https://yourcompany.atlassian.net`.
- d. **Authentication** – Choose to **Create and add a new secret** or **Use an existing one**.

If you choose to create a new secret, a Secrets Manager secret window opens requesting the following information:

- i. **Secret name** – A name for your Secrets Manager secret.
 - ii. **Jira username** – The username for your Jira user.
 - iii. **Jira password/API token** – The password/API token for your Jira user.
7. **Tags – optional** – Add an optional tag to track your plugin.
 8. Choose **Save**.

AWS CLI

To create a Jira plugin

```
aws qbusiness create-plugin \  
--application-id application-id \  
--display-name display-name \  
--type JIRA \  
--server-url https://example.atlassian.net \  
--auth-configuration basicAuthConfiguration="{secretArn=<secret-arn>,roleArn=<role-arn>}"
```

Configuring a Salesforce plugin

Salesforce is a customer relationship management (CRM) tool for managing support, sales, and marketing teams that you can use to create cases (tickets) to track issues. If you're a Salesforce user, you can create an Amazon Q Business plugin to allow your end users to create Salesforce cases from within their web experience chat.

To create a Salesforce plugin, you need configuration information from your Salesforce instance to set up a connection between Amazon Q and Salesforce and allow Amazon Q to perform actions in Salesforce.

For more information on how to use plugins during your web experience chat, see [Using plugins](#).

Topics

- [Prerequisites](#)
- [Service access roles](#)
- [Creating a plugin](#)

Prerequisites

Before you configure your Amazon Q Salesforce plugin, you must do the following:

- Set up a Connected App using the admin role in your Salesforce instance with Client Credentials Flow enabled.
- As an admin, configure an execution user with scoped permissions for performing actions in Amazon Q. For instructions, see [Configure a Connected App for the OAuth 2.0 Client Credentials Flow](#) in the Salesforce documentation.
- Note your Salesforce Connected App's consumer key (`client_id`) and your Salesforce Connected App Consumer secret (`client_secret`). You will need this OAuth 2.0 authentication information for creating an AWS Secrets Manager secret during the plugin configuration process.
- Note the Salesforce My Domain URL of your Salesforce organization. For example: `https://yourdomain.my.salesforce.com`.

Service access roles

To successfully connect Amazon Q to Salesforce, you need to give Amazon Q the following permission to access your Secrets Manager secret to get your Salesforce credentials. Amazon Q assumes this role to access your Salesforce credentials.

The following is the service access IAM role required:

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
```

```
        "Action": [
            "secretsmanager:GetSecretValue"
        ],
        "Resource": [
            "arn:aws:secretsmanager:{{your-region}}:{{your-account-id}}:secret:
[[secret-id]]"
        ]
    }
]
```

If you use the console and choose to create a new IAM role, Amazon Q creates the role for you. If you use the console and choose to use an existing secret, or you use the API, make sure your IAM role contains these permissions.

Creating a plugin

To create a Salesforce plugin for your web experience chat, you can use the AWS Management Console or the [CreatePlugin](#) API operation. The following tabs provide a procedure for creating a Salesforce plugin using the console and code examples for the AWS CLI.

Console

To create a Salesforce plugin

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. From the Amazon Q console, in **Applications**, select the name of your application from the list of applications.
3. From the left navigation menu, choose **Enhancements**, and then choose **Plugins**.
4. For **Plugins**, choose **Add plugin**.
5. For **Add plugins**, choose **Salesforce**.
6. For **Salesforce**, enter the following information:
 - a. **Name**, for **Plugin name** – A name for your Amazon Q plugin. The name can include hyphens (-), but not spaces, and can have a maximum of 1,000 alphanumeric characters.
 - b. **Service access** – Choose **Create and add a new service role** or **Use an existing service role**. Make sure that your service role has the necessary permissions.

- c. **URL** – My Domain URL of your Salesforce organization. For example: `https://yourdomain.my.salesforce.com`
- d. **Authentication** – Choose **Create and add a new secret** or **Use an existing one**. Your secret must contain the following information:
 - i. **Secret name** – A name for your Secrets Manager secret.
 - ii. **Connected app consumer key** – The consumer key for your Salesforce connected app.
 - iii. **Connected app consumer secret** – The consumer secret for your Salesforce connected app.
7. **Tags** – *optional* – An optional tag to track your plugin.
8. Choose **Save**.

AWS CLI

To create a Salesforce plugin

```
aws qbusiness create-plugin \  
--application-id application-id \  
--display-name display-name \  
--type SALESFORCE \  
--server-url //example.my.salesforce.com \  
--auth-configuration oAuth2ClientCredentialConfiguration="{secretArn=<secret-  
arn>,roleArn=<role-arn>}"
```

Configuring a ServiceNow plugin

ServiceNow provides a cloud-based service management system to create and manage organization-level workflows, such as IT services, ticketing systems, and support. ServiceNow uses incidents (tickets) to track issues. If you're a ServiceNow user, you can create an Amazon Q Business plugin to allow your end users to create ServiceNow cases from within their web experience chat.

To create a ServiceNow plugin, you need configuration information from your ServiceNow instance to set up a connection between Amazon Q and ServiceNow and allow Amazon Q to perform actions in ServiceNow.

For more information on how to use plugins during your web experience chat, see [Using plugins](#).

Topics

- [Prerequisites](#)
- [Service access roles](#)
- [Creating a plugin](#)

Prerequisites

Before you configure your Amazon Q ServiceNow plugin, you must do the following:

- As an admin, set up a new user in your ServiceNow instance with scoped permissions for performing actions in Amazon Q.
- Note your ServiceNow username and ServiceNow password. You will need this basic authentication information for creating an AWS Secrets Manager secret during the plugin configuration process.
- Note the base URL of your ServiceNow instance. For example: `https://yourinstance.service-now.com`.

Service access roles

To successfully connect Amazon Q to ServiceNow, you need to give Amazon Q the following permission to access your Secrets Manager secret to get your ServiceNow credentials. Amazon Q assumes this role to access your ServiceNow credentials.

The following is the service access IAM role required:

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Action": [
      "secretsmanager:GetSecretValue"
    ],
    "Resource": [
      "arn:aws:secretsmanager:{{your-region}}:{{your-account-id}}:secret:
[[secret-id]]"
    ]
  }]
}
```

```
    }  
  ]  
}
```

If you use the console and choose to create a new IAM role, Amazon Q creates the role for you. If you use the console and choose to use an existing secret, or you use the API, make sure your IAM role contains these permissions.

Creating a plugin

To create a ServiceNow plugin for your web experience chat, you can use the AWS Management Console or the [CreatePlugin](#) API operation. The following tabs provide a procedure for creating a ServiceNow plugin using the console and code examples for the AWS CLI.

Console

To create a ServiceNow plugin

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. From the Amazon Q console, in **Applications**, select the name of your application from the list of applications.
3. From the left navigation menu, choose **Enhancements**, and then choose **Plugins**.
4. For **Plugins**, choose **Add plugin**.
5. For **Add plugins**, choose **ServiceNow**.
6. For **ServiceNow**, enter the following information:
 - a. **Name**, for **Plugin name** – A name for your Amazon Q plugin. The name can include hyphens (-), but not spaces, and can have a maximum of 1,000 alphanumeric characters.
 - b. **Service access** – Choose **Create and add a new service role** or **Use an existing service role**. Make sure that your service role has the necessary permissions.
 - c. **URL** – The base URL of your ServiceNow instance. For example: `https://yourinstance.service-now.com`
 - d. **Authentication** – Choose **Create and add a new secret** or **Use an existing one**. Your secret must contain the following information:
 - i. **Secret name** – A name for your Secrets Manager secret.

- ii. **ServiceNow username** – The username for your ServiceNow user.
 - iii. **ServiceNow password** – The password for your ServiceNow user.
7. **Tags** – *optional* – An optional tag to track your plugin.
 8. Choose **Save**.

AWS CLI

To create a ServiceNow plugin

```
aws qbusiness create-plugin \  
--application-id application-id \  
--display-name display-name \  
--type SERVICE-NOW \  
--server-url //example.service-now.com \  
--auth-configuration basicAuthConfiguration="{secretArn=<secret-arn>,roleArn=<role-arn>}"
```

Configuring a Zendesk plugin

Zendesk is a customer relationship management system that helps businesses automate and enhance customer support interactions by creating tickets to track work. If you're a Zendesk user, you can create an Amazon Q Business plugin to allow your end users to create Zendesk cases from within their web experience chat.

To create a Zendesk plugin, you need configuration information from your Zendesk instance to set up a connection between Amazon Q and Zendesk and allow Amazon Q to perform actions in Zendesk.

For more information on how to use plugins during your web experience chat, see [Using plugins](#).

Topics

- [Prerequisites](#)
- [Service access roles](#)
- [Creating a plugin](#)

Prerequisites

Before you configure your Amazon Q Zendesk plugin, you must do the following:

- As an admin, set up a new user in your Zendesk instance with scoped permissions for performing actions in Amazon Q.
- (Optional) [Create an API token](#) for that new user.
- Note your Zendesk username and Zendesk password/API token. You will need this basic authentication information for creating an AWS Secrets Manager secret during the plugin configuration process.
- Note the base URL of your Zendesk instance. For example: `https://yoursubdomain.zendesk.com`.

Service access roles

To successfully connect Amazon Q to Zendesk, you need to give Amazon Q the following permission to access your Secrets Manager secret to get your Zendesk credentials. Amazon Q assumes this role to access your Zendesk credentials.

The following is the service access IAM role required:

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Action": [
      "secretsmanager:GetSecretValue"
    ],
    "Resource": [
      "arn:aws:secretsmanager:{{your-region}}:{{your-account-id}}:secret:
[[secret-id]]"
    ]
  }
]
```

If you use the console and choose to create a new IAM role, Amazon Q creates the role for you. If you use the console and choose to use an existing secret, or you use the API, make sure your IAM role contains these permissions.

Creating a plugin

To create a Zendesk plugin for your web experience chat, you can use AWS Management Console or the [CreatePlugin](#) API operation. The following tabs provide a procedure for creating a Zendesk plugin using the console and code examples for the AWS CLI.

Console

To create a Zendesk plugin

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. From the Amazon Q console, in **Applications**, select the name of your application from the list of applications.
3. From the left navigation menu, choose **Enhancements**, and then choose **Plugins**.
4. For **Plugins**, choose **Add plugin**.
5. For **Add plugins**, choose **Zendesk**.
6. For **Zendesk**, enter the following information:
 - a. **Name, Plugin name** – A name for your Amazon Q plugin. The name can include hyphens (-), but not spaces, and can have a maximum of 1,000 alphanumeric characters.
 - b. For **Service access** – Choose **Create and add a new service role** or **Use an existing service role**. Make sure that your service role has the necessary permissions.
 - c. **URL** – The base URL of your Zendesk instance. For example: `https://yoursubdomain.zendesk.com`
 - d. **Authentication** – Choose **Create and add a new secret** or **Use an existing one**. Your secret must contain the following information:
 - i. **Secret name** – A name for your Secrets Manager secret.
 - ii. **Zendesk username** – The username for your Zendesk user.
 - iii. **Zendesk password/API token** – The password/API token for your Zendesk user.
7. **Tags** – *optional* – An optional tag to track your plugin.
8. Choose **Save**.

AWS CLI

To create a Zendesk plugin

```
aws qbusiness create-plugin \  
--application-id application-id \  
--display-name display-name \  
--type ZENDESK \  
--server-url //example.zendesk.com \  
--auth-configuration basicAuthConfiguration="{secretArn=<secret-arn>,roleArn=<role-arn>}"
```

Using Amazon Q Business plugins

After plugins have been configured, you can use them to perform supported actions in your Amazon Q Business web experience chat. This topic provides an overview of how to use plugins.

Important

Once configured, all authorized Amazon Q web experience end users can use plugins to perform supported actions. If a plugin is activated for an application, end users will see an option to **Use a plugin**. If a plugin is deactivated, users won't see an option to use a plugin. End user access to plugins can't be customized.

Topics

- [Performing a plugin action](#)
- [Example plugin action prompts](#)

Performing a plugin action

The following describes how to perform a plugin action from within a web experience chat using both the console and the API.

Console

Performing a plugin action

1. Navigate to the deployed web experience URL and sign with your credentials on the login screen.
2. From conversation settings, choose **Use a plugin**.
3. You can choose to enact plugin actions in two ways:
 - a. Ask to perform an action directly. For example: Create a Jira ticket for a broken mouse. See [Quick create](#) for more details.
 - b. Start chatting in your web experience to find answers to your questions. Then choose to include the conversation context in any plugin action that you take. For example: Summarize this conversation and create a Jira ticket. For more information, see [Contextual create](#).
4. In response to your prompt for an action, Amazon Q displays a review form where you fill in the necessary information required to successfully complete an action.
5. To successfully complete the action, you need to submit it. Your web experience will display a success message if the action succeeds, or an error message if the action fails.

API

Performing a plugin action

```
aws qbusiness --no-verify-ssl --endpoint-url $endpoint \  
chat-sync --application-id application-id --user-id user-id \  
--user-message "Create an issue in Jira for broken button in web application" --  
chat-mode PLUGIN_MODE \  
--chat-mode-configuration '{  
  "pluginConfiguration": {  
    "pluginId": "plugin-id"  
  }  
}'
```

Example plugin action prompts

There are two ways you can choose to use plugins in your web experience chat, *quick creation* and *contextual creation*.

Topics

- [Quick create](#)

- [Contextual create](#)

Quick create

Using quick creation you can directly instruct your web experience to perform a plugin action. For example:

- Create a Zendesk ticket for a broken mouse
- Log an incident in ServiceNow for network outage
- Cut an issue in Jira for a broken link on a web page
- Create a Salesforce case for a missing invoice

Contextual create

Using contextual creation you can include conversation contexts to create tickets. For example, consider the following example conversation flows:

Example contextual create actions

- [Example 1: Create a ServiceNow incident](#)
- [Example 2: Create a ZenDesk ticket](#)
- [Example 3: Create a Salesforce case](#)
- [Example 4: Create a Jira issue](#)

Example 1: Create a ServiceNow incident

- **User prompt 1** – How to resolve network issues
- **Amazon Q response** – *Sample response*
- **User prompt 2** – How to reset my router
- **Amazon Q response** – *Sample response*
- **User action request** – Summarize this conversation and create a ServiceNow incident

Example 2: Create a ZenDesk ticket

- **User prompt 1** – Compare Amazon Kendra with OpenSearch

- **Amazon Q response** – *Sample response*
- **User action request** – Create a Zendesk ticket to migrate to Amazon Kendra

Example 3: Create a Salesforce case

- **User prompt 1** – Where is the IT office located
- **Amazon Q response** – *Sample response*
- **User prompt 2** – What floor is the office located in
- **Amazon Q response** – *Sample response*
- **User action request** – Create a case in Salesforce summarizing this conversation

Example 4: Create a Jira issue

- **User prompt 1** – How do I enable auto-scaling in EC2
- **Amazon Q response** – *Sample response*
- **User prompt 2** – How do I create an auto-scaling group
- **Amazon Q response** – *Sample response*
- **User action request** – Summarize this conversation and create an issue in Jira

Managing Amazon Q Business plugins

To manage Amazon Q plugins, you can take the following actions:

Actions

- [Updating a plugin](#)
- [Deleting a plugin](#)
- [Getting plugin properties](#)
- [Listing plugins](#)

Updating a plugin

To update a plugin, you can use AWS Management Console or the [UpdatePlugin](#) API operation. The following tabs provide a procedure for the console and code examples for the AWS CLI.

Console

To update a plugin

1. Sign in to the AWS Management Console and open the Amazon Q console.
2. From the Amazon Q console, in **Applications**, select the name of your application from the list of applications.
3. From the left navigation menu, choose **Enhancements**, and then choose **Plugins**.
4. For **Plugins**, select the plugin that you want to update, and then choose **Actions**.
5. For **Actions**, choose **Edit**.

On the plugins configuration page, you can edit your settings.

To deactivate a plugin

1. Sign in to the AWS Management Console and open the Amazon Q console.
2. From the Amazon Q console, in **Applications**, select the name of your application from the list of applications.
3. From the left navigation menu, choose **Enhancements**, and then choose **Plugins**.
4. For **Plugins**, select the plugin that you want to deactivate, and then choose **Actions**.
5. For **Actions**, choose **Deactivate**.

Your plugin will be deactivated. After your plugin is deactivated, its status will change to **Inactive**.

To reactivate a plugin

1. Sign in to the AWS Management Console and open the Amazon Q console.
2. From the Amazon Q console, in **Applications**, select the name of your application from the list of applications.
3. From the left navigation menu, choose **Enhancements**, and then choose **Plugins**.
4. For **Plugins**, select the plugin that you want to reactivate, and then choose **Actions**.
5. For **Actions**, choose **Reactivate**.

Your plugin will be activated. After your plugin is reactivated, its status will change to **Active**.

AWS CLI

To edit a plugin

```
aws qbusiness update-plugin \  
--application-id application-id \  
--plugin-id plugin-id \  
--display-name display-name \  
--server-url https://example.atlassian.net \  
--auth-configuration basicAuthConfiguration="{secretArn=<secret-arn>,roleArn=<role-arn>}"
```

To disable a plugin

```
aws qbusiness update-plugin \  
--application-id application-id \  
--plugin-id plugin-id \  
--state DISABLED
```

To enable a plugin

```
aws qbusiness update-plugin \  
--application-id application-id \  
--plugin-id plugin-id \  
--state ENABLED
```

Deleting a plugin

To delete a plugin, you can use the AWS Management Console or the [DeletePlugin](#) API operation. The following tabs provide a procedure for the console and code examples for the AWS CLI.

Console

To delete a plugin

1. Sign in to the AWS Management Console and open the Amazon Q console.
2. From the Amazon Q console, in **Applications**, select the name of your application from the list of applications.

3. From the left navigation menu, choose **Enhancements**, and then choose **Plugins**.
4. For **Plugins**, select the plugin that you want to delete, and then choose **Actions**.
5. For **Actions**, choose **Delete**.
6. In the dialog box, type **delete** to confirm your action.

The console displays a successful deletion message when the plugin deletion process is finished.

AWS CLI

To delete a plugin

```
aws qbusiness delete-plugin \  
--application-id application-id \  
--plugin-id plugin-id
```

Getting plugin properties

To get the details of an Amazon Q plugin, you can use either the AWS Management Console or the [GetPlugin](#) API operation. The following tabs provide a procedure for the console and code examples for the AWS CLI.

Console

To get plugin details

1. Sign in to the AWS Management Console and open the Amazon Q console.
2. From the Amazon Q console, in **Applications**, select the name of your application from the list of applications.
3. From the left navigation menu, choose **Enhancements**, and then choose **Plugins**.
4. For **Plugins**, select the configured plugin that you want to see details for.
5. On the **Plugin settings** page, the following details are available:
 - **Name** – The name of your plugin.
 - **Type** – The type of your plugin.

- **AWS Secrets Manager** – The Secrets Manager secret.
- **Creation time** – The time stamp for when your plugin was created.
- **Plugin ID** – The ID that's assigned to your plugin.

AWS CLI

To get plugin details

```
aws qbusiness get-plugin \  
--application-id application-id \  
--plugin-id plugin-id
```

Listing plugins

To list Amazon Q plugins, you can use the AWS Management Console or the [ListPlugins](#) API operation. The following tabs provide a procedure for the console and code examples for the AWS CLI.

Console

To list plugins

1. Sign in to the AWS Management Console and open the Amazon Q console.
2. From the Amazon Q console, in **Applications**, select the name of your application from the list of applications.
3. From the left navigation menu, choose **Enhancements**, and then choose **Plugins**.
4. In **Plugins**, a list of plugins that are attached to your application is available.

AWS CLI

To list plugins

```
aws qbusiness list-plugins \  
--application-id application-id
```

Admin controls and guardrails in Amazon Q Business

With Amazon Q Business, you can customize your application to your organizational needs. Amazon Q offers application *guardrails* or *chat controls* that you can configure to control the end user chat experience.

Using the guardrails feature, you can define global controls and topic-level controls for your application like the following:

- Control whether end users can upload files in chat to generate responses from uploaded files.
- Specify whether all Amazon Q chat responses will be generated using only enterprise data or whether your application can also use its underlying large language model (LLM) to generate responses when it can't find answers in your enterprise data.
- Control how Amazon Q responds to specific topics in chat.
- Customize which users and groups Amazon Q topic-level controls apply to.

Topics

- [Key terms for Amazon Q Business guardrails and chat controls](#)
- [Using global controls in Amazon Q Business](#)
- [Using topic-level controls in Amazon Q Business](#)
- [Managing Amazon Q Business admin controls and guardrails](#)

Key terms for Amazon Q Business guardrails and chat controls

The following are key terms you should know to understand guardrails in Amazon Q Business:

- **Enterprise data** – Data connected to your application using either an Amazon Q connector, direct document upload, or through an Amazon Kendra retriever.
- **Model knowledge** – The underlying knowledge outside your enterprise data that your large language model (LLM) is trained on.
- **Topic** – An admin user defined natural language topic.
- **Global controls** – Application level controls for controlling the sources that your application uses to generate responses (model knowledge and enterprise data, or enterprise data only). Global controls also define and control blocked phrases within your application.

- **Topic controls** – Topic-specific controls to determine the web application's behavior when it encounters a mention of a blocked topic by an end user.
- **Rules** – An application behavior logic configured to manage a controlled topic for a particular group of users.

Using global controls in Amazon Q Business

You can use Amazon Q Business global controls to configure settings that apply to conversations in your application.

Note

You can't create or delete guardrail global controls. You can only update existing global controls in your application.

The following are the global features that you can customize:

Global controls

- [Response settings](#)
- [Blocked phrases](#)
- [Feature control](#)
- [Customizing global controls](#)

Response settings

By default, a Amazon Q Business application is configured to respond to end user chat queries using only enterprise data. If it can't find information from your data sources, it responds with: "Sorry, I couldn't find enough information to answer."

When you update your application guardrails, you can use **Response settings** to change this behavior in the following ways:

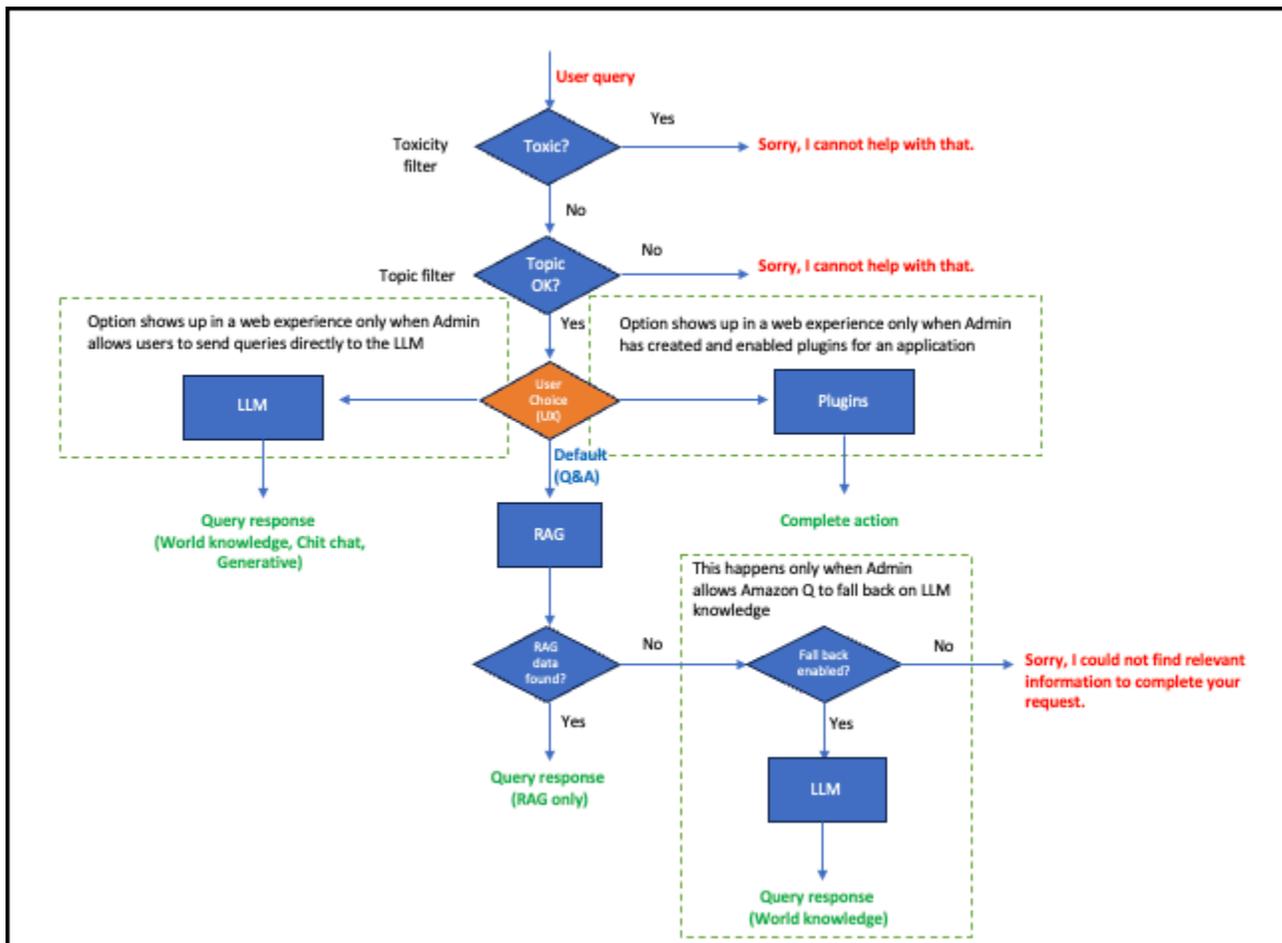
- **Allow end users to send queries directly to the LLM** – Give end users the option to either generate LLM-only responses or only generate responses from connected data sources. If you choose to activate this option, end users will be able to toggle between generating responses

from either the data sources you have connected to your application or use only the LLM to generate responses.

If you choose to activate this feature for your end users, they will see the option to turn **All data sources off** or **Respond from approved sources** in their web experience. If you turn the this feature off, then this option won't be available—or displayed—to end users in a web experience.

- **Allow Amazon Q to fall back to LLM knowledge** – Allow Amazon Q to use its LLM knowledge to generate responses when it can't find responses from your connected data sources. If you choose to activate this mode, and haven't given your end users the option to choose how responses are generated, your application will default to producing responses using the LLM when it can't find information in your data sources.

The following diagram shows you how Amazon Q uses these guardrails to direct queries:



⚠ Important

If you're changing response settings for an Amazon Q application created and deployed before 16 April, 2024, you need to update your web experience service role. For information on service role permissions needed, see [IAM role for an Amazon Q web experience](#). For information on how to update your web experience service role, see [Updating a web experience](#).

ℹ Note

Displaying sample prompts to your end user using the Amazon Q [Quick prompts](#) feature might not work if you choose to restrict response generation to enterprise data.

Global controls apply to all supported conversation interactions, except when it conflicts with a specific topic-level control. In that case, a topic-level control takes precedence.

Blocked phrases

You can define blocked phrases for your application. Amazon Q ensures that chat responses don't include these words. You can choose up to 20 words.

Additionally, you can optionally configure a custom message to be displayed to your end users in response to any mention of blocked phrases during chat. You can use this message to inform them that word is blocked and provide them with further guidance on next steps.

By default, your application doesn't define any blocked words. You can choose to add these words when you edit and update your global control guardrails.

Feature control

You can control whether end users can upload files during chat to ask questions based on the uploaded document.

By default, your application allows your end users to directly upload files in chat.

Customizing global controls

When you create an Amazon Q application, it's assigned the following default global controls:

- Generate responses from enterprise data only.
- No blocked words allowed.
- File upload by end user during chat is activated.

To update global topic controls for your web experience chat, you can use the AWS Management Console or the [UpdateChatControlsConfiguration](#) API operation. The following tabs provide a procedure for the console and code examples for the AWS CLI.

Note

You can't create or delete guardrail global controls. You can only update existing global controls in your application.

Console

To update a global control guardrail

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, select the name of your application from the list of applications.
3. From the left navigation menu, choose **Enhancements**, and then choose **Guardrails**.
4. In **Guardrails**, from **Global controls**, choose **Edit**.
5. In **Application guardrails**, do the following:
 - For **Response settings** do the following:
 - **Allow end users to send queries directly to the LLM** – If you choose to activate this option, end users will be able to toggle between generating responses from either the data sources you have connected to your application or use only the LLM to generate responses.

Note

If you choose to enable this option, your end users will have the option to generate LLM-only responses even if you don't allow Amazon Q to use LLM knowledge to generate responses.

For more information, see [Using global controls in Amazon Q](#).

- **Allow Amazon Q to fall back to LLM knowledge** – Choose this option if you want to generate responses from your application's LLM world knowledge when it can't find information in your connected data sources. The default is to restrict responses to enterprise data. For more information, see [Using global controls in Amazon Q](#).
 - For **Blocked words** – Define blocked words for the application. The application will not respond to questions that contain these words or mention them in any responses.
 - For **Messaging shown for blocked words** – Choose to create a custom response for your end users informing them of blocked word usage and any next steps to take.
6. For **Feature settings**, choose whether your end users will be allowed to upload files directly in chat to ask questions based on file content.
 7. Choose **Save**.

AWS CLI

To update a global control guardrail

```
aws qbusiness update-chat-controls-configuration \  
--application-id application-id \  
--blocked-phrases-configuration-update '{"blockedPhrasesToCreateOrUpdate":["example phrase 1", "example phrase 2"],"blockedPhrasesToDelete":["example phrase 1", "example phrase 2"],"systemMessageOverride":"user facing message when blocked phrase encountered"}' \  
--client-token clientToken \  
--response-scope ENTERPRISE_CONTENT_ONLY | EXTENDED_KNOWLEDGE_ENABLED
```

Using topic-level controls in Amazon Q Business

You can use topic-level controls to specify special topics within your application. You can configure rules to customize how Amazon Q Business should respond when a chat message matches a special topic. To streamline your application's response, you provide a name and a short description for how the large language model (LLM) should respond based on the topic-specific guardrail you're building. You can configure up to 10 topic-level controls.

Topic-level controls provide fine-grained customization for your application. For example, you can define a global control guardrail that allows your application to generate responses using model knowledge. You can also use a content retrieval rule to limit response generation for specific topics to enterprise content.

The following are the topic-level guardrails that you can customize:

Topic level guardrails

- [LLM prompt control](#)
- [Application behavior rules](#)
- [Creating topic controls](#)

LLM prompt control

You can add up to 5 representative messages that you expect end users to submit about this topic. You can also configure natural language descriptions to define the boundaries of the topic. Amazon Q Business uses these messages to check the responses that it generates for restricted content.

Application behavior rules

You can configure behavior rules that control how Amazon Q responds for each special topic that you specify.

Note

You can specify up to 5 rules per special topic.

Rules

- [Answer using enterprise data](#)
- [Blocking special topics](#)

Answer using enterprise data

When your application encounters a special topic, you can choose to allow it to answer from your enterprise data. If you allow responses from your enterprise data, you can further restrict which data sources in your application that your responses are generated from.

You can also choose to specify the specific users or groups within your application to apply this rule to, using either an inclusion logic or an exclusion logic. You can't use both kinds of logic at once. If a user is a member of a group with conflicting rules defined, Amazon Q will apply the more restrictive rule to that user.

Blocking special topics

When your application encounters a special topic, you can choose to block responses completely. If you do so, you can configure a custom message to display to your end users in response to any mention of blocked words during chat. Use this message to inform your end users that the topic is blocked and provide them with further guidance on next steps.

You can also choose to specify the specific groups within your application to apply this rule to, using either an inclusion logic or an exclusion logic. You can't use both kinds of logic at once. If a user is a member of a group with conflicting rules defined, Amazon Q will apply the more restrictive rule to that user.

Not specifying an inclusion or exclusion logic will result in the rule being applied to all users.

Creating topic controls

To create an Amazon Q topic-level control for your web experience chat, you can use AWS Management Console or the [UpdateChatControlConfiguration](#) operation. The following tabs provide a procedure for the console and code examples for the AWS CLI.

Console

To create a topic control

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, select the name of your application from the list of applications.
3. From the left navigation menu, choose **Enhancements**, and then choose **Guardrails**.
4. For **Guardrails**, from **Topic specific controls**, choose **Create topic control**.
5. For **Create topic specific controls**, enter the following information:
 - **Name** – Enter a name for your topic-specific control.
 - **Description** – A natural language description for your topic control configuration. Use this to help the LLM better identify queries associated with the topic control you're configuring.

6. For **Example chat messages**, enter representative phrases that you expect a user to type to invoke this topic. You can add up to 5 messages.
7. (Optional) To configure a rule, choose **Add new rule**.
8. For **Rule 1**, enter the following information:
 - In **Behavior in response to guardrail**, for **Behavior** – Choose how Amazon Q will respond to blocked topics: **Answer using enterprise data** or **Block completely**.
 - If you choose **Block completely** – Choose to include a custom message to inform your end user of restricted topics from chat and suggest follow up actions.
 - If you choose **Answer using enterprise data**, **Data source requirements** – Choose data sources that Amazon Q will use to generate responses.
9. For **User handling**, specify the users or groups that this topic control rule applies to and any users or groups that are exempt from this rule.
10. Choose **Save**.

AWS CLI

To create a topic control

```
aws qbusiness update-chat-controls-configuration \
--application-id application-id \
--client-token clientToken \
--topic-configurations-to-create-or-update
' [{"name": "name", "description": "description", "exampleChatMessages":
[ "message1", "message2" ], "rules": [ { "includedUsersAndGroups": { "userIds":
[ "userId1", "userId2" ], "userGroups": [ "userGroup1", "userGroup2" ] }, "ruleType":
"CONTENT_BLOCKER_RULE", "ruleConfiguration": { "contentBlockerRule":
{ "systemMessageOverride": "custom_message" } } }, { "excludedUsersAndGroups":
{ "userIds": [ "id1", "id2" ], "userGroups": [ "group1", "group2" ] }, "ruleType":
"CONTENT_RETRIEVAL_RULE", "ruleConfiguration": { "contentRetrievalRule":
{ "eligibleDataSources": [ { "indexId": "index-id1", "dataSourceId": "data-source-id1" },
{ "indexId": "index-id2", "dataSourceId": "data-source-id2" } ] } } } ] } ] ' \
--topic-configurations-to-delete ' { "name": "existing-topic-name" } '
```

Note

The user IDs you add to configure topic controls must already exist in your Identity Provider (IdP). You are responsible for validating any user groups you add.

Managing Amazon Q Business admin controls and guardrails

To manage Amazon Q Business admin controls and guardrails, you can take the following actions:

Note

You can't create or delete guardrail global controls. You can only update existing global controls in your application.

Actions

- [Deleting topic controls](#)
- [Getting topic control properties](#)

Deleting topic controls

To delete configured chat controls, you can use AWS Management Console or the [DeleteChatControlsConfiguration](#) API operation. The following tabs provide a procedure for the console and code examples for the AWS CLI.

Console

To delete topic controls

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, select the name of your application from the list of applications.
3. From the left navigation menu, choose **Enhancements**, and then choose **Guardrails**.
4. In **Guardrails**, from **Topic specific controls**, choose the topic control you want to delete, and then choose **Delete**.

5. In the dialog box, type **delete** to confirm your action.

The console displays a successful deletion message when the plugin deletion process is finished.

AWS CLI

To delete a topic specific control

```
aws qbusiness delete-chat-controls-configuration \  
--application-id application-id
```

Getting topic control properties

To get the details of Amazon Q topic controls, you can use either the AWS Management Console or the [GetChatControlsConfiguration](#) API operation. The following tabs provide a procedure for the console and code examples for the AWS CLI.

Console

To get configured details for admin controls and guardrails

1. Sign in to the AWS Management Console and open the Amazon Q console.
2. From the Amazon Q console, in **Applications**, select the name of your application from the list of applications.
3. From the left navigation menu, choose **Enhancements**, and then choose **Admin controls and guardrails**.

You will find the details of your configured **Global controls** and **Topic specific controls** on the page.

AWS CLI

To get admin controls and guardrails details

```
aws qbusiness get-chat-control-configuration \  

```



```
--application-id application-id
```

Boosting chat responses using relevance tuning

Note

This section assumes that you understand [document attributes and how they work](#) in Amazon Q.

If you choose to use an [Amazon Q native retriever](#), you can assign weights to document attributes after mapping them to Amazon Q index fields using the Amazon Q *relevance tuning* feature. Then, you can use these assigned weights to fine-tune the underlying ranking of RAG retrieved passages within your Amazon Q application to optimize the relevance of chat responses. In Amazon Q, boosting means to raise a document in chat results using these weights.

Important

Boosting document attributes using *relevance tuning* is an admin-only feature.

Boosting chat responses based on document attributes helps you rank sources that are more authoritative higher than other sources in your application. You can assign a higher value to more recent content, specific file types, or specific data sources.

Amazon Q automatically boosts specific document attributes, like document title, when retrieving information from your index to generate end user chat responses. You can use the boosting feature to customize and control boosting, and also override any pre-existing boosts applied by Amazon Q.

When you use this feature, a Retrieval Augmented Generation (RAG)-generated result is given a boost in the chat response when the query includes terms that match that field or attribute. You specify how much of a boost the document receives when there is a match. When Amazon Q generates responses, it prioritizes the sources that are assigned higher rankings.

Choosing to boost document attributes doesn't by itself cause Amazon Q to include or exclude a document in the chat response. A boosted document attribute is only one of the factors that Amazon Q uses to determine the relevance of a document.

Note

Boosting in Amazon Q is only available if you use an Amazon Q native retriever. If you use an Amazon Kendra retriever, you must [configure boosting for document attributes](#) in Amazon Kendra. Amazon Q supports any boosting that's already configured in your Amazon Kendra index.

Topics

- [Understanding boosting](#)
- [Boosting types](#)
- [Configuring document attributes for boosting](#)
- [Enabling document attributes for search](#)

Understanding boosting

To improve retrieved results and customize the end user chat experience, Amazon Q enables you to map attributes to fields in your Amazon Q index.

Amazon Q offers [two kinds of attributes](#):

- **Reserved or default** – Reserved attributes are based on document attributes that commonly occur in most data. You can use reserved attributes to map commonly occurring document attributes in your data to Amazon Q index fields.
- **Custom** – You can create custom attributes to map document attributes that are unique to your data to Amazon Q index fields.

Document attributes can be [mapped to index fields](#) using either the Amazon Q console or the API:

- **Use the API** – Before you use the API, you must first create an index. Next, create index fields. Then, to ingest documents into your Amazon Q index, use the [CreateDataSource](#) or [BatchPutDocument](#) API operations.
- **Use the console** – You can choose to map document attributes from your data sources when you connect your data source to Amazon Q. When you use the console, Amazon Q automatically maps data source document fields to Amazon Q index fields internally.

Document attributes—both reserved and custom—can only be of the following data types: DATE, NUMBER, STRING, and STRING_LIST. To use STRING and STRING_LIST type document attributes for boosting on the console and the API, they must be enabled for search. To enable these attributes, use the [DocumentAttributeConfiguration](#) object of the [UpdateIndex](#) API operation. If you don't enable search on these attributes, you can't boost attributes of these data types on either the Amazon Q console or the API.

To customize and control boosting for document attributes, use the `boostingOverride` parameter of the [NativeIndexConfiguration](#) object of the [UpdateRetriever](#) API operation.

For more information about Amazon Q document attributes and how to map them, see [Document attributes and types](#).

Boosting types

Amazon Q offers two types of boosting: document attribute boosting and document attribute value boosting. This section outlines how these types of boosting work.

Note

To use the STRING and STRING_LIST type document attributes for boosting on the console and the API, they must be enabled for search using the [DocumentAttributeConfiguration](#) object of the [UpdateIndex](#) API operation. If you don't enable search on these attributes, you can't boost attributes of these data types on either the Amazon Q console or the API.

Types of boosting

- [Boosting document attribute importance](#)
- [Boosting document attribute value](#)

Boosting document attribute importance

You can boost document attributes to control the relative importance, or boosting level, of the field for end user queries. You can boost importance for all document attribute data types that are supported by Amazon Q—DATE, NUMBER, STRING, and STRING_LIST.

Note

Amazon Q automatically boosts the document title attribute to **Low**. You can change this value when you customize boosting.

If you choose to boost document attributes, you can also customize boosting in the following ways:

- **Boost duration** – Specifies the time period over which a boost applies to a DATE type document attribute. For example, if you set boosting duration to 604,800 seconds (1 week) for the `_created_at` reserved attribute, documents created within the last week will be boosted.

Generally, all documents inside the boosting duration will be given more importance over documents outside the boosting duration. Within the boosting duration, documents with more recent dates will be given more importance over documents with less recent dates.

Outside the boosting duration, the documents with more recent dates will continue to be given more importance over documents with less recent dates. However, the overall effect of the date boosting will taper to zero as the dates move further away from the boosting duration.

Note

Boosting duration is based on the most recent date in all documents in the index.

- **Boost order** – Determines whether a NUMBER type document attribute is boosted in prioritizing higher values or prioritizing lower values.

For example, if your documents contain attributes for view count, you can choose to prioritize chat responses with higher view count values by boosting larger values over smaller values. Or, suppose your documents contain attributes that denote priority—for example, a task tracker that assigns priority 1 to the most important task. In that case, you can choose to boost documents using smaller values.

Boosting document attribute value

To customize boosting levels, you can boost document attribute values for only STRING type document attributes.

For example, suppose that you're applying an importance boost to a STRING attribute called department. The department attribute has values like HR and Legal. You can assign the values HR, VERY_HIGH and Legal, HIGH to customize the importance that Amazon Q gives to these attribute values when they match a chat request.

Configuring document attributes for boosting

To boost specific documents for end user queries using document attributes, you can use the AWS Management Console or the [DocumentAttributeBoostingConfiguration](#) parameter of the [UpdateRetriever](#) API operation.

How you configure relevance tuning depends on whether you integrate your Amazon Q application with IAM Identity Center (IDC) as an [AWS-managed](#) application, or use an external identity provider (IDP) for your Amazon Q application (including using IAM Identity Center as an IdP by creating a customer-managed IAM Identity Center app). The following sections provide procedures to boost document attributes using the console and code examples for the AWS CLI for both use cases.

Note

For STRING and STRING_LIST type document attributes to be used for boosting on the console and the API, they must be enabled for search using the [DocumentAttributeConfiguration](#) object of the [UpdateIndex](#) API operation. If you don't enable search on these attributes, you can't boost attributes of these data types on either the console or the API.

Topics

- [Configuring boosting for an IAM Identity Center powered app](#)
- [Configuring boosting for an external IdP powered app](#)

Configuring boosting for an IAM Identity Center powered app

The following tabs provide a procedure to boost document attributes using the console and code examples for the AWS CLI if you're using IAM Identity Center as an IdP.

Console

To boost document attributes

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, select the name of your application from the list of applications.
3. From the left navigation menu, choose **Relevance tuning**.
4. In **Relevance tuning**, choose the document attribute type that you want to boost.

 **Note**

You can boost attributes using the following values: **None**, **Low**, **Medium**, **High**, and **Very high**.

Choose from the following options:

- a. **Popular** – Amazon Q displays the following popularly boosted document attributes for you to choose from:
 - i. **Document title** – Use to boost the title of a document. You can also use **Advanced settings** to boost specific document titles. By default, the document title attribute is enabled for search with a value of Low. You can change this value when you customize boosting.
 - ii. **Last updated** – Use to boost content by its last updated date. You can also use **Advanced settings** to configure **Boosting duration**, or how long your boost should apply.
 - iii. **File type** – Use to boost content by file type.
 - iv. **Data sources** – Use to boost the content data source type.
 - v. To save your configuration, choose **Save**.
- b. **Text** – Use to boost STRING and STRING_LIST type reserved or custom document attributes that you have enabled for search. Then, choose **Save**.
- c. **Date** – Use to boost content using DATE type reserved or custom document attributes. For example, use the **Created at** document attribute to boost content based on recency. You can also use **Advanced settings** to configure **Boosting duration**, or how long your boost should apply. Then, choose **Save**.
- d. **Numeric** – Use to boost content using NUMERIC type reserved or custom attributes. For example, use the **View count** document attribute to boost content based on

view count. Based on your boosting needs, choose either **Prioritize higher values** or **Prioritize lower values**. Then, choose **Save**.

- e. Once done, you can select **View web experience** to check boosting. Your configured web experience will open in a new window.

AWS CLI

Update your Amazon Q index to apply boosting

This example shows how to apply VERY_HIGH boosting for the STRING type document attribute `_document_title`.

```
aws qbusiness update-retriever \
--application-id APPLICATION-ID --retriever-id RETRIEVER-ID \
--configuration '{
  "nativeIndexConfiguration": {
    "indexId": "INDEX-ID",
    "boostingOverride": {
      "_document_title": {
        "stringConfiguration": {
          "boostingLevel": "VERY_HIGH"
        }
      }
    }
  }
}'
```

This example shows how to apply boosting for the STRING type attribute `_category`, the DATE type attribute `_created_at`, the NUMBER type attribute `_view_count`, and the STRING_LIST type attribute `_authors`.

```
aws qbusiness update-retriever \
--application-id APPLICATION-ID --retriever-id RETRIEVER-ID \
--configuration '{
  "nativeIndexConfiguration": {
    "indexId": "INDEX-ID",
    "boostingOverride": {
      "_category": {
        "stringConfiguration": {
```

```

        "boostingLevel": "LOW",
        "attributeValueBoosting": {
            "HR": "MEDIUM"
        }
    },
    "_created_at": {
        "dateConfiguration": {
            "boostingLevel": "LOW",
            "boostingDurationInSeconds": 2592000
        }
    },
    "_view_count": {
        "numberConfiguration": {
            "boostingLevel": "LOW",
            "boostingType": "PRIORITIZE_SMALLER_VALUES"
        }
    },
    "_authors": {
        "stringListConfiguration": {
            "boostingLevel": "HIGH"
        }
    }
}
}'

```

Update your Amazon Q retriever to remove any existing boosts

This example shows how to remove any existing boosts from document attributes in your retriever.

```

aws qbusiness update-retriever \
--application-id APPLICATION-ID --retriever-id RETRIEVER-ID \
--configuration '{
    "nativeIndexConfiguration": {
        "indexId": "INDEX-ID"
    }
}'

```

Get details about your Amazon Q retriever boosts

This example shows how to get details for your existing boosting configuration

```
aws qbusiness get-retriever \  
--application-id APPLICATION-ID --retriever-id RETRIEVER-ID
```

Configuring boosting for an external IdP powered app

The following tabs provide a procedure to boost document attributes using the console and code examples for the AWS CLI if you're using an external IdP.

Console

To boost document attributes

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, select the name of your application from the list of applications, and then choose **Preview web experience**.
3. In **Preview web experience**, from the edit menu on the right, choose **Boost**.
4. From the **Boosting** menu, choose the document attribute type that you want to boost.

Note

You can boost attributes using the following values: **None**, **Low**, **Medium**, **High**, and **Very high**.

Choose from the following options:

- a. **Popular** – Amazon Q displays the following popularly boosted document attributes for you to choose from:
 - i. **Document title** – Use to boost the title of a document. You can also use **Advanced settings** to boost specific document titles. By default, the document title attribute is enabled for search with a value of Low. You can change this value when you customize boosting.

- ii. **Last updated** – Use to boost content by its last updated date. You can also use **Advanced settings** to configure **Boosting duration**, or how long your boost should apply.
 - iii. **File type** – Use to boost content by file type.
 - iv. **Data sources** – Use to boost the content data source type.
 - v. To save your configuration, choose **Save**.
- b. **Text** – Use to boost STRING and STRING_LIST type reserved or custom document attributes that you have enabled for search. Then, choose **Save**.
 - c. **Date** – Use to boost content using DATE type reserved or custom document attributes. For example, use the **Created at** document attribute to boost content based on recency. You can also use **Advanced settings** to configure **Boosting duration**, or how long your boost should apply. Then, choose **Save**.
 - d. **Numeric** – Use to boost content using NUMERIC type reserved or custom attributes. For example, use the **View count** document attribute to boost content based on view count. Based on your boosting needs, choose either **Prioritize higher values** or **Prioritize lower values**. Then, choose **Save**.

AWS CLI

Update your Amazon Q index to apply boosting

This example shows how to apply VERY_HIGH boosting for the STRING type document attribute `_document_title`.

```
aws qbusiness update-retriever \  
--application-id APPLICATION-ID --retriever-id RETRIEVER-ID \  
--configuration '{  
    "nativeIndexConfiguration": {  
        "indexId": "INDEX-ID",  
        "boostingOverride": {  
            "_document_title": {  
                "stringConfiguration": {  
                    "boostingLevel": "VERY_HIGH"  
                }  
            }  
        }  
    }  
}
```

```
}'
```

This example shows how to apply boosting for the STRING type attribute `_category`, the DATE type attribute `_created_at`, the NUMBER type attribute `_view_count`, and the STRING_LIST type attribute `_authors`.

```
aws qbusiness update-retriever \
--application-id APPLICATION-ID --retriever-id RETRIEVER-ID \
--configuration '{
  "nativeIndexConfiguration": {
    "indexId": "INDEX-ID",
    "boostingOverride": {
      "_category": {
        "stringConfiguration": {
          "boostingLevel": "LOW",
          "attributeValueBoosting": {
            "HR": "MEDIUM"
          }
        }
      },
      "_created_at": {
        "dateConfiguration": {
          "boostingLevel": "LOW",
          "boostingDurationInSeconds": 2592000
        }
      },
      "_view_count": {
        "numberConfiguration": {
          "boostingLevel": "LOW",
          "boostingType": "PRIORITIZE_SMALLER_VALUES"
        }
      },
      "_authors": {
        "stringListConfiguration": {
          "boostingLevel": "HIGH"
        }
      }
    }
  }
}'
```

Update your Amazon Q retriever to remove any existing boosts

This example shows how to remove any existing boosts from document attributes in your retriever.

```
aws qbusiness update-retriever \  
--application-id APPLICATION-ID --retriever-id RETRIEVER-ID \  
--configuration '{  
    "nativeIndexConfiguration": {  
        "indexId": "INDEX-ID"  
    }  
}'
```

Get details about your Amazon Q retriever boosts

This example shows how to get details for your existing boosting configuration

```
aws qbusiness get-retriever \  
--application-id APPLICATION-ID --retriever-id RETRIEVER-ID
```

Enabling document attributes for search

For STRING and STRING_LIST type attributes to be eligible for boosting, they must first be enabled for search in your Amazon Q index. To enable these attributes for search, use the [DocumentAttributeConfiguration](#) object of the [UpdateIndex](#) API operation.

The following sections provide AWS CLI examples of how to enable document attributes for search.

Topics

- [Making reserved document attributes searchable](#)
- [Making custom document attributes searchable](#)
- [Checking document attribute search activation](#)

Making reserved document attributes searchable

The following is an example of how to use the AWS CLI to enable for search the STRING type reserved document attribute `_category` and the STRING_LIST type reserved document attribute `_authors` by using the [UpdateIndex](#) API operation.

```
aws qbusiness update-index \  
--application-id APPLICATION_ID \  
--index-id INDEX_ID \  
--document-attribute-configurations '  
  [  
    {  
      "name": "_category",  
      "type": "STRING",  
      "search": "ENABLED"  
    },  
    {  
      "name": "_authors",  
      "type": "STRING_LIST",  
      "search": "ENABLED"  
    }  
  ]'
```

Making custom document attributes searchable

You can also enable custom document attributes for search using the [DocumentAttributeConfiguration](#) object of the [UpdateIndex](#) API operation.

The following is an example of how to use the AWS CLI to enable for search the custom STRING and STRING_LIST type document attributes using the [UpdateIndex](#) API operation.

```
aws qbusiness update-index \  
--application-id APPLICATION_ID \  
--index-id INDEX_ID \  
--document-attribute-configurations '  
  [  
    {  
      "name": "custom_string",  
      "type": "STRING",  
      "search": "ENABLED"  
    },  
    {  
      "name": "custom_string_list",  
      "type": "STRING_LIST",  
      "search": "ENABLED"  
    }  
  ]'
```

Checking document attribute search activation

To check if a `STRING` or `STRING_LIST` type document attribute has been enabled for search successfully, use the [GetIndex](#) API operation.

```
aws qbusiness get-index \  
--application-id APPLICATION_ID \  
--index-id INDEX_ID
```

The AWS CLI returns the following type of response:

```
{  
  ...  
  "documentAttributeConfigurations": [  
    {  
      "name": "_authors",  
      "search": "ENABLED",  
      "type": "STRING_LIST"  
    },  
    {  
      "name": "_category",  
      "search": "ENABLED",  
      "type": "STRING"  
    },  
    {  
      "name": "_created_at",  
      "search": "DISABLED",  
      "type": "DATE"  
    },  
    {  
      "name": "_data_source_id",  
      "search": "ENABLED",  
      "type": "STRING"  
    },  
    {  
      "name": "_document_title",  
      "search": "ENABLED",  
      "type": "STRING"  
    },  
    {  
      "name": "_file_type",  
      "search": "ENABLED",
```

```
    "type": "STRING"
  },
  {
    "name": "_language_code",
    "search": "ENABLED",
    "type": "STRING"
  },
  {
    "name": "_last_updated_at",
    "search": "DISABLED",
    "type": "DATE"
  },
  {
    "name": "_source_uri",
    "search": "ENABLED",
    "type": "STRING"
  },
  {
    "name": "_version",
    "search": "ENABLED",
    "type": "STRING"
  },
  {
    "name": "_view_count",
    "search": "DISABLED",
    "type": "NUMBER"
  }
],
...
}
```

Customizing an Amazon Q web experience

If you're using IAM Identity Center to manage user access to your application, you can customize the Amazon Q web experience that you created for your end users in the AWS console. You do this after you create and enhance an Amazon Q application. When you customize your web experience, you can personalize it by changing its title and subtitle adding a welcome message, and displaying sample prompts.

Note

You can't run any chat queries from the web experience customize mode.

You can customize a web experience by using either the AWS Management Console or the Amazon Q API. If you use the API, customizing your Amazon Q can involve a combination of the following API operations:

- [CreateApplication](#) – Creates an Amazon Q application
- [CreateWebExperience](#) – Creates an Amazon Q web experience
- [GetWebExperience](#) – Gets the properties of the web experience that you set up
- [ListWebExperiences](#) – Lists Amazon Q web experiences that you created

If you use the console to create your Amazon Q application, a web experience is created automatically and connected to your chosen data source. You can customize that web experience on the **Customize web experience** console page.

Before you can customize a web experience, you must complete [creating your application](#).

Topics

- [Customize web experience](#)
- [Managing Amazon Q web experiences](#)

Customize web experience

The following tabs provide a procedure for customizing a web experience on the AWS Management Console and code examples for the AWS CLI.

Console

To customize an Amazon Q web experience

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps to [selecting an Amazon Q retriever](#), [connecting data sources](#) and [adding users and groups](#).
4. Then, from the Amazon Q application page, select **Customize web experience**.
5. In **Customize web experience**, from the right navigation pane, select **Customize web experience**.
6. In **Customize web experience**, enter the following information for your web experience:
 - **Title** – A title for your web experience. End users see this title on their web experience page.
 - **Subtitle - *optional*** – A subtitle for your web experience to highlight other information for your end users. This subtitle is visible to your end users on their web experience page.
 - **Welcome message** – Provide an optional welcome message for your end users. We recommend mentioning data sources and application capabilities.
 - **Display sample prompts** – Provide a list of [sample prompts](#) on the end user's conversation start screen.
7. Choose **Save**.
8. To exit the web experience customize and return to the Amazon Q console control panel to deploy your application, select **Sign out** from the left pane.

AWS CLI

To create and customize a web experience

```
aws qbusiness create-web-experience \  
--application-id application-id \  
--role-arn roleArn \  
--title optional-title \  
--subtitle optional-subtitle \  
--welcome-message optional-welcome-message \  

```

```
--sample-prompts-control-mode ENABLED
```

Managing Amazon Q web experiences

To manage Amazon Q web experiences, you can take the following actions:

Actions

- [Deleting a web experience](#)
- [Getting properties of a web experience](#)
- [Listing web experiences](#)

Deleting a web experience

To delete an Amazon Q web experience, you can use the console or the [DeleteWebExperience](#) API operation.

If you're using the API, you can delete a web experience without deleting the application that it's a part of.

If you're using the console, the only way to delete your Amazon Q web experience is to delete the Amazon Q application that it's attached to.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To delete an Amazon Q web experience

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, choose **Actions**.
3. Choose **Delete**.
4. In the dialog box that opens, type **Delete** to confirm deletion, and then choose **Delete**.

You are returned to the service console while your application is deleted. When the deletion process is complete, the console displays a message confirming successful deletion. Both the application and the web experience are deleted.

AWS CLI

To delete an Amazon Q web experience

```
aws qbusiness delete-web-experience \  
--application-id application-id \  
--web-experience-id web-experience-id
```

Getting properties of a web experience

To get the properties of an Amazon Q web experience, you can use the console or the [GetWebExperience](#) API operation.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To get properties of an Amazon Q web experience

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, select the name of your application from the list of applications.
3. For **Web experience settings**, the following settings are available:
 - **Web experience IAM role ARN** – The IAM role assumed by end users when they log in to your web experience.
 - **Deployed URL** – The deployed URL of your web experience.
 - **Tags** – Tags that are attached to your web experience.

To update a setting, choose **Edit**.

AWS CLI

To get properties of an Amazon Q web experience

```
aws qbusiness get-web-experience \  
--application-id application-id \  
--web-experience-id web-experience-id
```

Listing web experiences

To list Amazon Q web experiences, you can use the console or the [ListWebExperiences](#) API operation.

If you use the console, you can only see the web experience that's attached to a single application.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To list Amazon Q web experiences

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. For **Applications**, the Amazon Q web experience attached to your application is shown.

AWS CLI

To list Amazon Q web experiences

```
aws qbusiness get-web-experience \  
--application-id application-id \  
--max-results max-results-to-return
```

Using an Amazon Q Business web experience

An Amazon Q Business web experience is an easy-to-use generative artificial intelligence (generative AI) assistant. You can use the Amazon Q web experience to ask questions and to accomplish your tasks. When you ask a question, the Amazon Q web experience analyzes the latest approved data collected from various data sources within your organization to generate a comprehensive response.

With an Amazon Q web experience, you can ask complex questions in plain language and get a detailed response. You can also use an Amazon Q web experience to perform tasks for you, such as draft an email message or create a Jira ticket.

The Amazon Q web experience provides you with the following capabilities:

Web experience features

- [Prompts](#)
- [Engage with contextual responses](#)
- [Analyze content](#)
- [Perform actions on your behalf](#)
- [Review source citations](#)
- [Upload files and chat](#)
- [Copy responses](#)
- [Provide feedback](#)
- [Conversation management](#)
- [Conversation settings](#)
- [Retrieve conversation IDs](#)

Prompts

The welcome page optionally provides example prompts to help you understand the types of questions and tasks that you can ask the Amazon Q web experience. This feature is provided depending on how the web experience is configured. If provided, use the sample prompts to formulate your own questions and tasks.

Engage with contextual responses

The Amazon Q web experience analyzes your questions and returns responses that use information from various data sources within your organization. You can continue with the conversation in the context of the active session or start a new conversation.

Analyze content

Ask the Amazon Q web experience to summarize its response, generate text from the response, do comparative analysis, and also perform math and reasoning tasks.

Perform actions on your behalf

Use the Amazon Q web experience to perform actions on your behalf using [plugins](#). For example, you can ask the web experience to schedule a meeting, create a ticket in Jira, or draft an email message. You only see an option to **Use a plugin** in your web experience if your admin has enabled it. You can only choose to perform plugin actions with **Use a plugin** mode enabled.

Review source citations

The Amazon Q web experience provides in-text source citations in the form of a numbered list. To view the source of the response, choose the number at the end of the sentence.

To view the entire list of sources, choose **Sources** at the end of the response. Use the source list to fact-check the response or for deeper analysis.

Upload files and chat

With the Amazon Q web experience, you can upload documents that aren't stored in your organization's data sources and knowledge base. Then you can use the uploaded documents to ask questions and summarize or analyze data that's based on the content of the uploaded documents. The uploaded documents aren't stored and are available for use only during the session in which the documents are uploaded.

To upload documents during a session, choose the upload icon next to the question box. You can upload a maximum of five files in a single session.

Copy responses

You can copy and save the responses for later review and analysis. To copy a response, choose the copy icon at the end of the response.

Provide feedback

To provide immediate feedback about the response you received from the Amazon Q web experience, use the thumbs-up or thumbs-down button. Your feedback is used to help address technical issues in the web experience.

If you select the thumbs-down button, you can choose from the following feedback options:

- **Response is not helpful (incorrect or not relevant to my query)**
- **Response is not based on company documents**
- **Response is not complete**
- **Response is not concise**
- **The sources are inaccurate or missing**
- **Other (explain below)**

You can add additional context for any of the thumbs-down feedback options you choose in the **Additional details (optional)** text box.

Conversation management

Amazon Q web experience stores each conversation for up to 30 days. Your conversations are listed in the left navigation pane. You can perform the following tasks to manage your conversations:

- **View conversation history** – Choose a conversation to view the conversation history for that session.
- **Start new conversation** – Choose **+ New conversation** to start a new conversation.
- **Delete conversation** – Choose a conversation that you want to delete, choose **Delete**, and then choose **Delete** again.

Conversation settings

If your admin has allowed you to, you can use choose to configure Amazon Q web experience responses in two ways from **Conversation settings**:

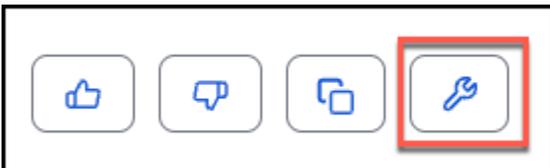
- **Respond from approved sources** – If you select this mode, Amazon Q will only choose to retrieve data from your enterprise to generate responses.
- **All data sources off** – If you select this mode, Amazon Q will choose to respond from your application's underlying world knowledge only.

For more information, see [Using global controls in Amazon Q](#).

Retrieve conversation IDs

To troubleshoot issues in your Amazon Q web experience, you may need your web experience conversation IDs. Provide these details in any support tickets you open.

If you are an admin user configuring a web experience using the Amazon Q console, you can use the spanner icon in your web experience preview to copy debugging information (application ID, conversation ID, user message ID, and system message ID) to share with support.



If you're an end user using a deployed web experience, you need to use developer tools in your browser to retrieve this information. The following tabs provide a procedure for retrieving this information from your web experience using different web browsers.

Chrome

To retrieve Amazon Q web experience conversation IDs

1. Make sure you're on the deployed web experience browser page you're having the conversation in.
2. Then, from the Google Chrome web browser menu, choose **More tools**, and then choose **Developer Tools**.

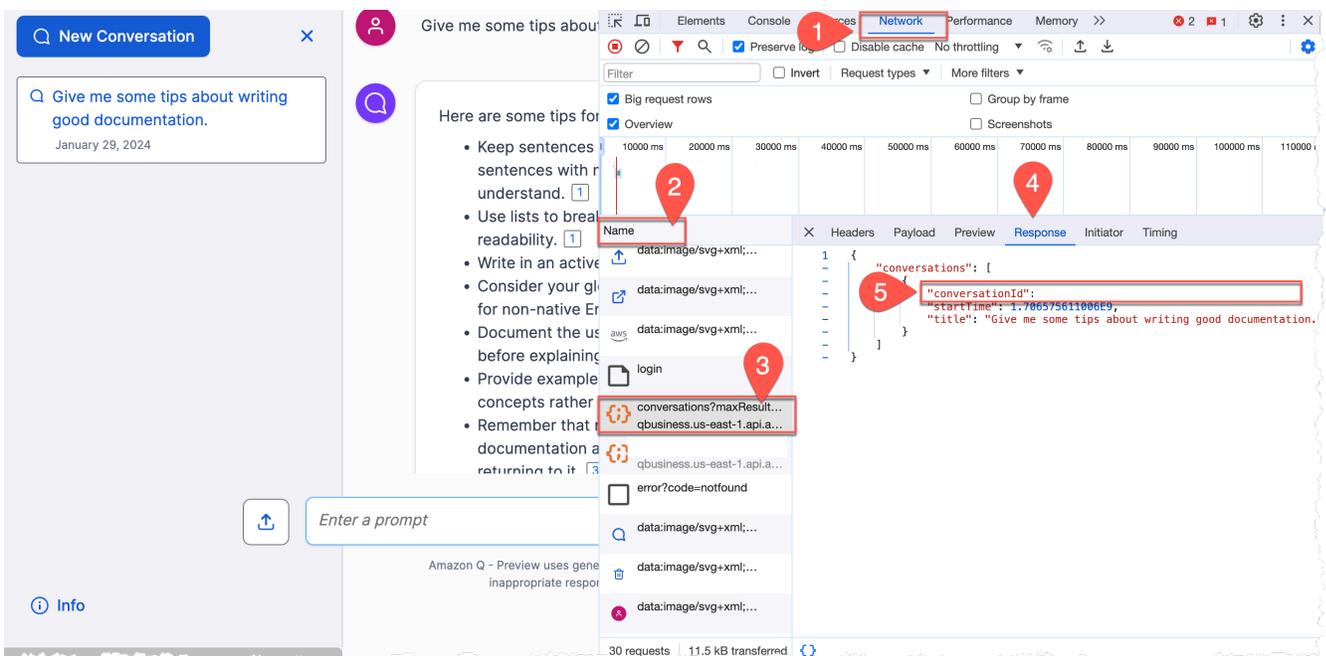
Note

You can also open the context (right-click) menu and then choose the **Inspect** option.

- Then, choose **Network**, and then choose the conversation from the **Name** column.
- Then, choose the **Response** tab. The response tab will show your conversation ID. You may need to re-select conversation from the Amazon Q left navigation menu to refresh **Network** activity and the inspection window.

Note

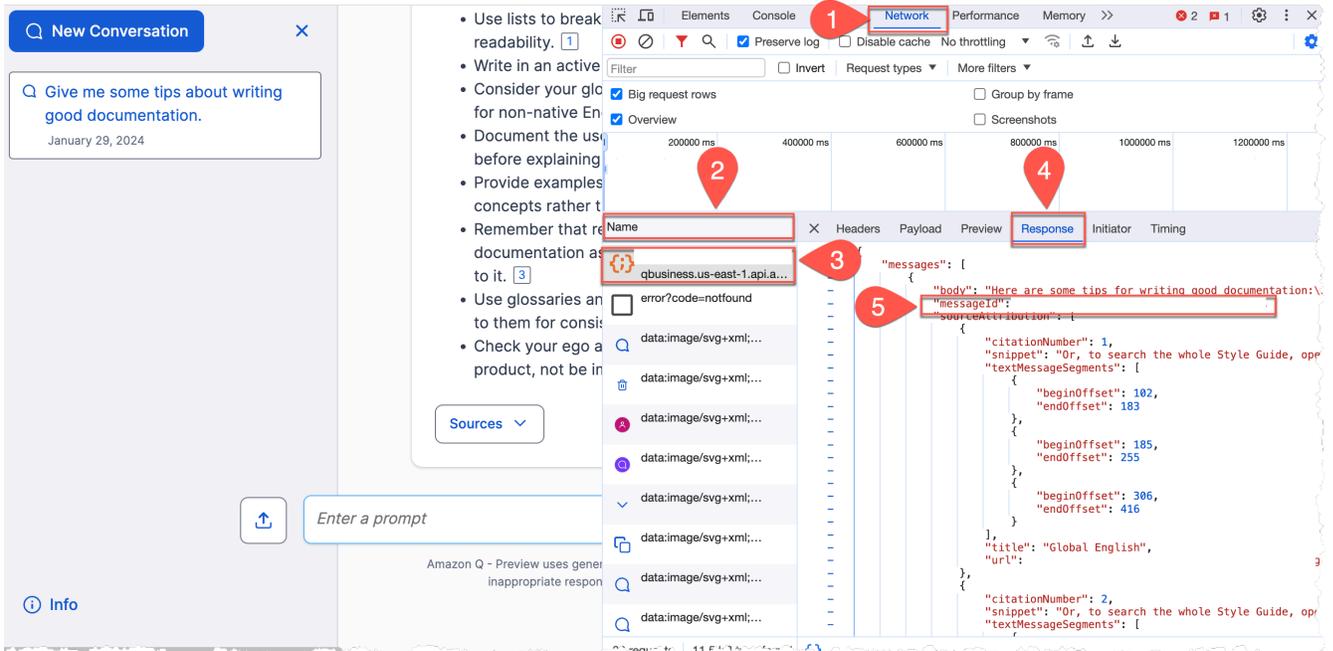
A collection of prompts and responses in a conversation (an entry in the left navigation menu of the web experience) is assigned a unique **Conversation ID**.



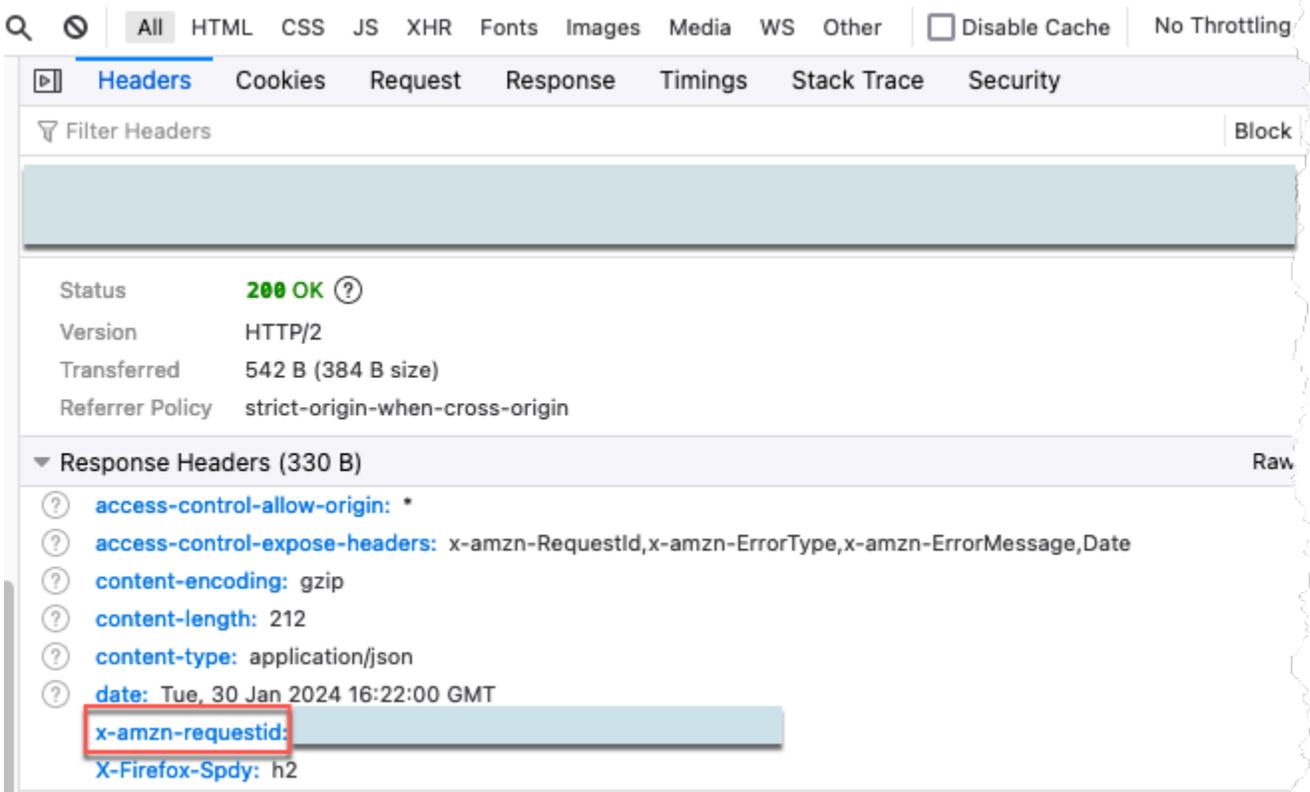
- Then, choose the message from the **Name** column and choose the **Response** tab. The response tab will show your message ID.

Note

The user typed prompt can be seen in the **Preview** tab and the Amazon Q response in the **Response** tab. Each tab has a **Message ID**.



- Then, choose the **Headers** tab, and go to **Response Headers**. The **Response Headers** tab will show your request id (**x-amzn-requestid**).



Firefox

To retrieve Amazon Q web experience conversation IDs

1. Make sure you're on the deployed web experience browser page you're having the conversation in.
2. Then, from the Firefox web browser menu, choose **More tools**, and then choose **Web Developer Tools**.

Note

You can also open the context (right-click) menu and then choose the **Inspect** option.

3. Then, choose **Network**, and then choose the conversation from the **Name** column.
4. Then, choose the **Response** tab. The response tab will show your conversation ID. You may need to re-select conversation from the Amazon Q left navigation menu to refresh **Network** activity and the inspection window

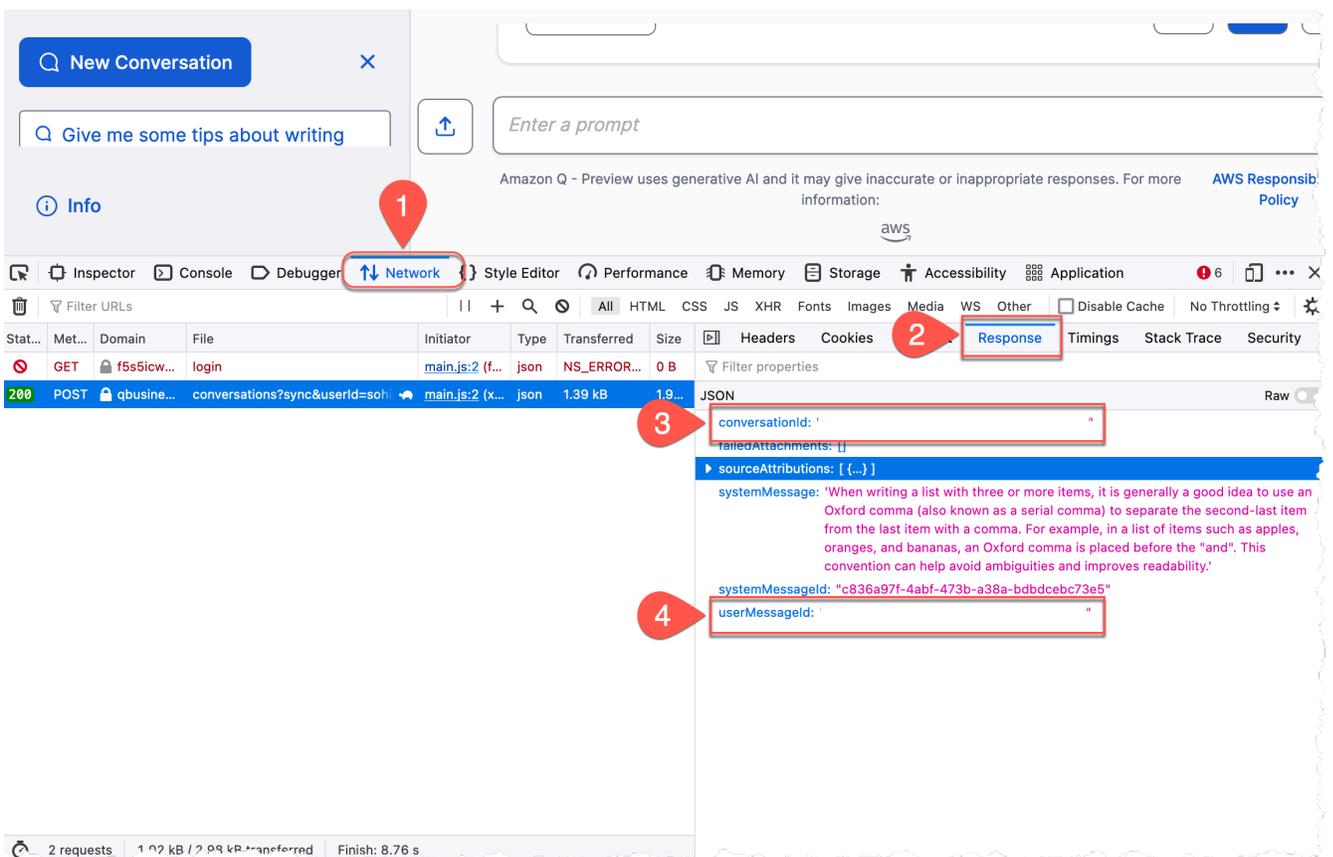
Note

A collection of prompts and responses in a conversation (an entry in the left navigation menu of the web experience) is assigned a unique **Conversation ID**.

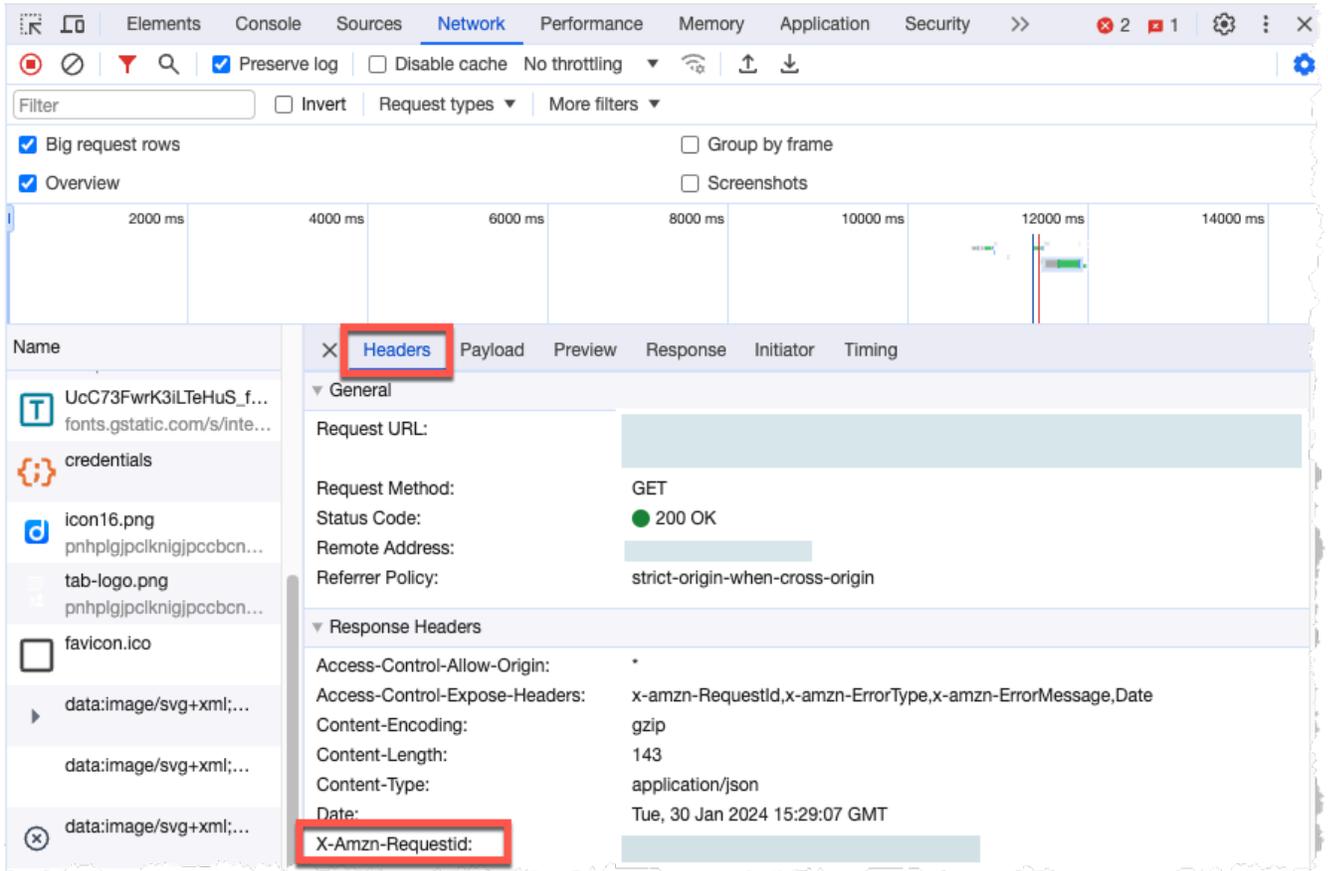
- Then choose the message from the **Name** column and choose the **Response** tab. The response tab will show your message ID.

Note

The user typed prompt can be seen in the **Preview** tab and the Amazon Q response in the **Response** tab. Each tab has a **Message ID**.



- Then, choose the **Headers** tab, and go to **Response Headers**. The **Response Headers** tab will show your request id (**X-Amzn-Requestid**).



Amazon Q Business features

In addition to [enhancements](#), Amazon Q Business offers the following features:

- **Filtering using metadata** – Use document attributes to customize and control the end user chat experience. Currently supported only if you use the Amazon Q API.
- **Source attribution with citations** – Verify responses using Amazon Q source attributions.
- **Upload files and chat** – Let end users upload files directly into chat and use uploaded file data to perform web experience tasks.
- **Quick prompts** – Feature sample prompts to inform end users of the capabilities of their Amazon Q web experience.

Topics

- [Filtering chat responses using document attributes](#)
- [Source attribution with citations in Amazon Q Business](#)
- [Upload files and chat in Amazon Q Business](#)
- [Quick prompts in Amazon Q Business](#)

Filtering chat responses using document attributes

Note

Prerequisite: This section assumes you have an understanding of [document attributes and how they work](#) in Amazon Q.

If you use the API, Amazon Q Business includes a filtering by document attribute feature. With this feature, you can customize and control chat responses for your end user using attributes—or metadata attached to documents mapped to index fields. For example, if data source type is an attribute attached to your documents, you can specify that chat responses be generated only from a specific data source.

Or, you can allow end users to restrict the scope of chat responses using the attribute filters that you have selected. For example, an end user can choose that their chat responses be generated using documents from specific data sources.

Filtering chat responses using metadata has the following key benefits:

- **Ensure response relevance and accuracy** – You can specify that responses be generated from and limited to authoritative sources within your data
- **Control response context** – You can specify the type (PDF, for example) and corpus (Business Requirement Documents, for example) of documents that responses will be generated from.
- **Maintain response freshness** – You can restrict chat responses to only documents that were generated after a specific date.
- **Scope chat responses** – You can help your end user narrow the scope of their responses and get to the right answer quicker.

Amazon Q offers a set of reserved document attributes that you can use. You can also create custom document attributes that are more representative of your organization's data and use cases for more fine-grained chat response control.

Important

During Preview, filtering using document attributes in chat is only supported through the API. Boosting search results using document attributes is supported on both the console and the API.

Source attribution with citations in Amazon Q Business

The Amazon Q Business web experience chat response provides in-text source citations for responses that use the organization's data sources and knowledge base as a source. The chat response also provides an entire list of sources used to generate the response.

In-text source citations

In-text citations are provided in the form of a numbered list at the end of a sentence. To view an in-text source citation, choose a citation number. Each citation provides the following attributes:

- **Title** – The title of the document that's the source for the generated response.
- **URL** – The URL of the document that's the source for the generated response. Choose the URL to view the source document.

Source list

Sources used to generate the response are provided at the end of the response. Each source listed provides the following attributes:

- **Citation number** – The number provided at the end of the sentences in the response.
- **Title** – The title of the document that's the source for the generated response.
- **Text segment** – A text extract from a source document that's used for source attribution.
- **URL** – The URL of the document that's the source for the generated response.

Upload files and chat in Amazon Q Business

End users using the Amazon Q Business web experience can upload documents that might not be stored in your organization's data sources and knowledge base. They can use the uploaded documents to ask questions and summarize or analyze data that's based on the content of the uploaded documents. The uploaded documents aren't stored and are available for use only for the conversation in which the documents are uploaded.

You can upload up to 5 files during a conversation. The size of each file you upload must be 10 MB or less. The total parsed content for all files combined have to be under 30,000 tokens or 20,000 words. 1 word corresponds roughly to 1.5 tokens.

Amazon Q supports specific document types for upload. To learn more about the document types that can be uploaded, see [Supported document formats in Amazon Q Business](#).

If you're uploading Comma Separated Values (CSV) or Microsoft Excel (XLS and XLSX) documents into chat, Amazon Q performs best for tables with approximately 4 columns and 10 rows.

Quick prompts in Amazon Q Business

The Amazon Q Business web experience welcome page provides sample prompts to help end users understand the types of questions and tasks that they can ask in the web experience. Sample prompts aren't enabled by default.

If you're an AWS Management Console customer and are configuring the web experience for your end users, you can enable the sample prompts feature when you preview the web experience. For more information, see [Preview and customize web experience](#).

⚠ Important

Before you enable the sample prompts feature, make sure that the **Only produce responses from retrieval augmented generation (RAG)** check box for the **Application guardrails** is cleared. For more information, see [Customizing global controls](#). The sample prompts might not work if the responses is restricted to enterprise data.

Currently, you can't create your own prompts or edit the provided sample prompts.

Using an external identity provider to manager user access

For applications using legacy identity management, Amazon Q Business requires that you integrate your web experience with an identity provider (IdP) that's compliant with SAML 2.0. This integration is required so that only authorized end users from within your organization have access to your content. Amazon Q Business can work with any IdP that's compliant with SAML 2.0. Amazon Q uses service-initiated single sign-on (SSO) to authenticate users. IdP-initiated SSO is not supported.

Starting April 30, 2024, all new applications will need to use IAM Identity Center directly to manage user access. All existing Amazon Q applications using [legacy identity management](#) will need to migrate to using IAM Identity Center for user management by July 31, 2024. We recommend you integrate any new application you're creating directly with IAM Identity Center.

Topics

- [Admin workflow using an external IDP](#)
- [Create an Amazon Q application](#)
- [Previewing and customizing an Amazon Q Business web experience](#)
- [Creating and selecting a retriever for an Amazon Q Business application](#)
- [Connecting data sources to an Amazon Q Business application](#)
- [Deploying an Amazon Q Business web experience](#)

Admin workflow using an external IDP

Important

Starting April 30, 2024, all new applications will need to use IAM Identity Center directly to manage user access. All existing Amazon Q applications using [legacy identity management](#) will need to migrate to using IAM Identity Center for user management by July 31, 2024. We recommend you integrate any new application you're creating directly with IAM Identity Center.

If you're an admin user using an external SAML 2.0 based identity provider (IDP) for your Amazon Q application (including using IAM Identity Center as a SAML 2.0 based IdP by creating a customer-managed IAM Identity Center app), you create and configure an Amazon Q web experience by completing the following steps:

1. [Creating the Amazon Q application](#) that powers your web experience.
2. [Choosing a retriever](#) for the application.
3. [Connecting your data sources](#) to—or uploading data into—the application.
4. [Enhancing and customizing the web experience](#) by configuring admin-level controls, and the end user chat experience. For more information, see [Enhancing an Amazon Q application](#) and [Amazon Q features](#).
5. [Previewing your web experience](#) to test how it looks and works for your end users. In this step, you add a title and subtitle for your web experience, and a welcome message for your end users. You can choose to chat in preview mode to test responses. Only public data with no access control is used to generate queries in preview mode.
6. [Deploying your web experience](#) for your end users by integrating with a SAML 2.0 supported identity provider (IdP). If you're using the console, this step involves switching between your IdP console and the Amazon Q console.

Note

During Preview, an Amazon Q application supports only 50 end users. If you need more capacity, contact [Support](#).

Create an Amazon Q application

Important

Starting April 30, 2024, all new applications will need to use IAM Identity Center directly to manage user access. All existing Amazon Q applications using [legacy identity management](#) will need to migrate to using IAM Identity Center for user management by July 31, 2024. We recommend you integrate any new application you're creating directly with IAM Identity Center.

As the first step towards creating an Amazon Q Business chat application for your end users, you configure an Amazon Q application. Then, you can optionally enhance it by customizing the end user experience. After this, you select and create a retriever, and connect and configure the data sources.

This section guides you through the process of creating and configuring an Amazon Q application. To create an application, you can use the Amazon Q console, the AWS Command Line Interface (AWS CLI), and the Amazon Q API operations.

As a prerequisite, make sure that you complete the [setting up](#) tasks. If you're using the AWS CLI or the API, make sure that you created the required [IAM roles](#).

After you finish creating your application, you can customize and preview the web experience that it will power.

Note

During Preview, an Amazon Q application supports only 50 end users. If you need more capacity, contact [Support](#).

The following tabs provide a procedure for creating an application that uses an external identity provider to manage user access. by using the AWS Management Console and code examples for using the AWS CLI.

Console

To configure an Amazon Q application

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. For **Create Amazon Q application**, choose **Get started**.
3. For **Applications**, choose **Create application**. The console will display a **Select access management method for application** dialog box.
4. In **Select access management method for application**, choose **Legacy identity management** and then select **Ok**. Choosing this option allows you to use SAML 2.0 to manage user identities using an identity provider of your choice.
5. For **Application settings**, enter the following information for your Amazon Q application:

- **Application name** – A name for your Amazon Q Business application for easy identification. This name is only visible in the AWS Management Console. The name can include hyphens (-), but not spaces, and can have a maximum of 1,000 alphanumeric characters.
- **Service access** – An IAM role for Amazon Q Business to allow it to access the AWS resources it needs to create your application. You can choose to use an existing role or create a new role.

 **Note**

For more information about example service roles, see [IAM role for an Amazon Q Business application](#).

- **Service role name** – A name for the service (IAM) role you created for easy identification on the console.
- **Encryption** – Amazon Q encrypts your data by default using AWS managed AWS KMS keys. To customize your encryption settings, select **Customize encryption settings (advanced)**. Then, you can choose to use an existing AWS KMS key or create a new one. To learn more, see [Data encryption](#).

 **Important**

If you choose to use a customer managed key, you must provision at least 10 index storage units when you [create an Amazon Q retriever](#).

6. **Tags – optional** – To add tags to your Amazon Q application and web experience, select **Add new tag**. Then, enter the following information for each tag:
 - **Key** – Add a key for your tag.
 - **Value - optional** – An optional value for your tag.

For more information about using tags with Amazon Q, see [Tags](#).

7. To start creating your application, choose **Create**.

AWS CLI

To configure an Amazon Q application

```
aws qbusiness create-application \  
--display-name application-name \  
--role-arn roleArn \  
--description application-description \  
--encryption-configuration kmsKeyId=<kms-key-id> \  
--attachments-configuration attachmentsControlMode=ENABLED
```

For information on managing your Amazon Q Business application, see [Managing applications](#).

Previewing and customizing an Amazon Q Business web experience

Important

Starting April 30, 2024, all new applications will need to use IAM Identity Center directly to manage user access. All existing Amazon Q applications using [legacy identity management](#) will need to migrate to using IAM Identity Center for user management by July 31, 2024. We recommend you integrate any new application you're creating directly with IAM Identity Center.

Important

A web experience preview is available only for existing applications using legacy identity management.

If you're integrating your Amazon Q application with an external SAML 2.0 compliant identity provider (IdP) (including using IAM Identity Center as your identity provider by creating a [customer managed](#) IAM Identity Center application), you can preview the Amazon Q web experience that you created for your end users in the AWS console. You do this after you create and enhance

an Amazon Q application. By previewing your web experience, you can test the features and enhancements that you configured for it.

Note

You can run a limited number of chat queries from the web experience preview. Only public documents ingested in your index are accessible—and used for generating responses—in the preview. Documents with access control are not accessible in, or searchable from, the preview.

You can customize and preview a web experience by using either the AWS Management Console or the Amazon Q API. If you use the API, previewing your Amazon Q can involve a combination of the following API operations:

- [CreateApplication](#) – Creates an Amazon Q application
- [CreateWebExperience](#) – Creates an Amazon Q web experience
- [GetWebExperience](#) – Gets the properties of the web experience that you set up
- [ListWebExperiences](#) – Lists Amazon Q web experiences that you created
- [ChatSync](#) – Starts or continues a conversation in your Amazon Q application

If you use the console to create your Amazon Q application, a web experience is created automatically and connected to your chosen data source. You can preview and deploy that web experience on the **Preview web experience** console page.

Before you can preview a web experience, make sure that you complete [creating your application](#).

Topics

- [Preview and customize web experience](#)
- [Testing Amazon Q Business web experience functions](#)
- [Managing Amazon Q Business web experiences](#)

Preview and customize web experience

The following tabs provide a procedure for previewing and customizing a web experience on the AWS Management Console and code examples for the AWS CLI.

Console

To preview and customize an Amazon Q web experience

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Optionally, complete the steps to [selecting an Amazon Q retriever](#), [connecting data sources](#), and [enhancing your application](#).
4. Then, from the Amazon Q application page, select **Preview web experience**.
5. In **Preview web experience**, from the right navigation pane, select **Customize web experience**.
6. In **Customize web experience**, enter the following information for your web experience:
 - **Title** – A title for your web experience. End users see this title on their web experience page.
 - **Subtitle - *optional*** – A subtitle for your web experience to highlight other information for your end users. This subtitle is visible to your end users on their web experience page.
 - **Display welcome message** – Provide an optional welcome message for your end users. We recommend mentioning data sources and application capabilities.
 - **Display sample prompts** – Provide a list of [sample prompts](#) on the end user's conversation start screen.
7. Choose **Save**.
8. To exit the web experience preview and return to the Amazon Q console control panel to deploy your application, select **Sign out** from the left pane.

AWS CLI

To create and customize a web experience

```
aws qbusiness create-web-experience \  
--application-id application-id \  
--title title \  
--subtitle subtitle \  
--welcome-message optional-welcome-message \  

```



```
--sample-prompts-control-mode ENABLED
```

Testing Amazon Q Business web experience functions

The following tabs provide a procedure for testing your web experience configuration for the AWS Management Console and code examples for the AWS CLI.

Console

To test your Amazon Q web experience chat

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Optionally, complete the steps to [selecting an Amazon Q retriever](#), [connecting data sources](#), and [enhancing your application](#).
4. Then, from the Amazon Q application page, select **Preview web experience**.
5. Choose from the following options to test your web experience:
 - a. **Ask questions** – Ask a question. Amazon Q generates and returns answers based on the enterprise data that the end user has access to. Continue the conversation by asking follow-up questions.
 - b. **Verify response sources** – Each Amazon Q answer cites the source documents used to generate it.
 - c. **See conversation history** – Amazon Q retains conversation history for 30 days so that they can search through questions and answers. You can view conversation history from the left navigation pane.
 - d. **Summarize content** – Amazon Q can summarize email message threads.
 - e. **Create outlines and drafts** – Use Amazon Q to create outlines and templates for documents.
 - f. **Perform plugin actions** – If you've configured [Plugins](#), ask Amazon Q to perform actions on your behalf, like creating a ticket in a supported third party app.
 - g. **Test guardrails and chat controls** – If you've configured [Guardrails and chat controls](#), check how Amazon Q responds to queries and special topics.

- To exit the web experience preview and return to the Amazon Q console control panel to deploy your application, select **Sign out** from the left pane.

AWS CLI

To preview web experience

```
aws qbusiness chat-sync \  
--application-id application-id \  
--user-id user-id \  
--user-groups user-groups \  
--user-message user message \  
--action-execution plugin-actions \  
--attachments file uploads \  
--attribute-filter attribute-filters
```

Managing Amazon Q Business web experiences

To manage Amazon Q Business web experiences, you can take the following actions:

Actions

- [Creating a web experience](#)
- [Deleting a web experience](#)
- [Getting properties of a web experience](#)
- [Listing web experiences](#)
- [Updating a web experience](#)

Creating a web experience

To create an Amazon Q web experience, you can use the console or the [CreateWebExperience](#) API operation.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

If you use the console, this action is spread across three steps: [Configuring an Amazon Q Business application](#), [Preview and customize web experience](#), and [Deploying an Amazon Q Business web experience](#). Amazon Q creates a web experience for you when you configure your application. To create a web experience, you must create an application.

AWS CLI

To create an Amazon Q web experience

```
aws qbusiness create-web-experience \
--application-id application-id \
--sample-prompts-control-mode sample-prompts \
--subtitle subtitle \
--tags tags \
--title title \
--welcome-message welcome-message \
```

Deleting a web experience

To delete an Amazon Q web experience, you can use the console or the [DeleteWebExperience](#) API operation.

If you're using the API, you can delete a web experience without deleting the application that it's a part of.

If you're using the console, the only way to delete your Amazon Q web experience is to delete the Amazon Q application that it's attached to.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To delete an Amazon Q web experience

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, choose **Actions**.

3. Choose **Delete**.
4. In the dialog box that opens, type **Delete** to confirm deletion, and then choose **Delete**.

You are returned to the service console while your application is deleted. When the deletion process is complete, the console displays a message confirming successful deletion. Both the application and the web experience are deleted.

AWS CLI

To delete an Amazon Q web experience

```
aws qbusiness delete-web-experience \  
--application-id application-id \  
--web-experience-id web-experience-id
```

Getting properties of a web experience

To get the properties of an Amazon Q web experience, you can use the console or the [GetWebExperience](#) API operation.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To get properties of an Amazon Q web experience

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, select the name of your application from the list of applications.
3. For **Web experience settings**, the following settings are available:
 - **Web experience IAM role ARN** – The IAM role assumed by end users when they log in to your web experience.
 - **Deployed URL** – The deployed URL of your web experience.
 - **Tags** – Tags that are attached to your web experience.

To update a setting, choose **Edit**.

AWS CLI

To get properties of an Amazon Q web experience

```
aws qbusiness get-web-experience \  
--application-id application-id \  
--web-experience-id web-experience-id
```

Listing web experiences

To list Amazon Q web experiences, you can use the console or the [ListWebExperiences](#) API operation.

If you use the console, you can only see the web experience that's attached to a single application.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To list Amazon Q web experiences

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. For **Applications**, the Amazon Q web experience attached to your application is shown.

AWS CLI

To list Amazon Q web experiences

```
aws qbusiness get-web-experience \  
--application-id application-id \  
--max-results max-results-to-return
```

Updating a web experience

To update an Amazon Q web experience, you can use the console or the [UpdateWebExperience](#) API operation.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To update an Amazon Q web experience

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, select the name of your application from the list of applications.
3. On your application page, select **Web experience settings** and then select **Edit**.
4. On the **Deploy web experience** page, you can edit your web experience settings.

AWS CLI

To update an Amazon Q web experience

```
aws qbusiness update-web-experience \  
--application-id application-id \  
--web-experience-id web-experience-id \  
--authentication-configuration authentication-configuration \  
--sample-prompts-control-mode sample-prompts \  
--subtitle subtitle \  
--title title \  
--welcome-message welcome-message
```

Creating and selecting a retriever for an Amazon Q Business application

Important

Starting April 30, 2024, all new applications using [legacy identity management](#) will need to use IAM Identity Center directly to manage user access. All existing Amazon Q applications will need to migrate to using IAM Identity Center for user management by July 31, 2024. We recommend you integrate any new application you're creating directly with IAM Identity Center.

After creating your Amazon Q Business application, you create and select the retriever that will power your generative AI web experience. A retriever pulls data from an index in real time during a conversation. Amazon Q provides retrievers for Amazon Kendra indexes and also for a native index. You can choose between selecting an Amazon Q retriever or using an already configured Amazon Kendra index as a retriever.

To select a retriever, you use the AWS Management Console or the [CreateRetriever](#) API operation.

If you use the console and choose to use a Amazon Q retriever, Amazon Q creates an index for you as part of the application configuration process. For easy tracking, you can tag both the retriever and index. If you use the API to create a Amazon Q retriever, you must also use the [CreateIndex](#) API operation to create an Amazon Q index.

Important

You can't change the retriever for your application after your application has been created. To change your retriever, you must create a new application.

Note

The data sources available to connect to your application change depending on your retriever choice.

For instructions on how to select a retriever, choose a topic based on your retriever preference for Amazon Q.

Topics

- [Creating an Amazon Q Business retriever](#)
- [Selecting an Amazon Kendra retriever to an Amazon Q Business application](#)

Creating an Amazon Q Business retriever

To select a Amazon Q Business retriever, you can use either the AWS Management Console, or the [CreateIndex](#) and [CreateRetriever](#) API operations.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To create an Amazon Q retriever

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Then, for **Select retriever**, choose **Use native retriever** – Build an Amazon Q Business retriever for your Amazon Q Business application. This option creates an Amazon Q Business index that can connect to the Amazon Q Business supported data sources that you choose.

Important

The native retriever includes a default capacity of 10k documents and 0.5 queries per second (QPS).

Note

Available data sources when you select this option include all [Amazon Q supported data connectors](#) and direct document upload.

4. For **Index provisioning** – Choose the **Number of units** that you need. Amazon Q charges you based on the document capacity that you choose. You can choose up to 50 units. Each unit is 20,000 documents or 200 MB, whichever comes first.
5. For **Tags** – Choose whether you want to add **Index tags**.
6. To create your retriever, choose **Create**.

AWS CLI

To create an Amazon Q index

```
aws qbusiness create-index \  
--application-id application-id \  
--display-name display-name \  
--description index-description \  
--capacity-configuration units =<index-capacity-units>
```

To create an Amazon Q retriever

```
aws qbusiness create-retriever \  
--application-id application-id \  
--display-name display-name \  
--type NATIVE_INDEX \  
--role-arn roleArn \  
--configuration nativeIndexConfiguration="{indexId=<created-index-id>}" \  
--tags tags
```

Managing Amazon Q Business retrievers

To manage Amazon Q Business retrievers, you can take the following actions:

Actions

- [Deleting an Amazon Q Business retriever](#)
- [Getting properties of an Amazon Q Business retriever](#)
- [Listing Amazon Q Business retrievers](#)
- [Updating Amazon Q Business retrievers](#)

Deleting an Amazon Q Business retriever

To delete a Amazon Q Business retriever and its associated index, you can use the console or the [DeleteRetriever](#) API operation.

If you use the `DeleteIndex` API operation, deleting a retriever also deletes the Amazon Q index that's attached to it. You can't selectively choose to delete an index attached to a retriever.

If you're using the console, the only way to delete your Amazon Q native retriever and the index associated with it, is to delete your Amazon Q application.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To delete an Amazon Q retriever

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, choose **Actions**.
3. Choose **Delete**.
4. In the dialog box that opens, type **Delete** to confirm deletion, and then choose **Delete**.

You are returned to the service console while your application is deleted. When the deletion process is complete, the console displays a message confirming successful deletion.

AWS CLI

To delete an Amazon Q retriever

```
aws qbusiness delete-retriever \  
--application-id application-id \  
--retriever-id retriever-id
```

Getting properties of an Amazon Q Business retriever

To get the properties of an Amazon Q Business retriever and index, you can use the console or the [GetRetriever](#) API operation.

Note

If you use the console, you can't edit or update retriever or index settings.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To get properties of an Amazon Q retriever

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, select the name of your application from the list of applications.
3. For **Retriever settings**, the following settings are available:
 - **Retriever** – The type of retriever that you're using.
 - **Document count** – The number of documents that are attached to your index.
 - **Last modified time** – The time that your index was last modified.
 - **Index ID** – The ID of the index attached to your retriever.
 - **Storage used** – The amount of storage that your index is using.
 - **Index status** – The status of your index.

AWS CLI

To get properties of an Amazon Q retriever

```
aws qbusiness get-retriever \  
--application-id application-id \  
--retriever-id retriever-id
```

Listing Amazon Q Business retrievers

To list your native Amazon Q Business retrievers, you can use the console or the [ListRetrievers](#) API operation.

If you use the console, the list of Amazon Q retrievers and indices attached to them correspond to the list of applications that you have created.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To list your Amazon Q retrievers

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. For **Applications**, a list of all retrievers (with indices associated) that you have created is available.

AWS CLI

To list your Amazon Q retrievers

```
aws qbusiness list-retrievers \  
--application-id application-id \  
--max-results maximum-result-to-display
```

Updating Amazon Q Business retrievers

To update your Amazon Q Business retriever, you can use the [UpdateRetriever](#) API operation.

You can't update your retriever and its associated index by using the console.

The following tab provides code examples for the AWS CLI.

Console

This action is not supported on the console.

AWS CLI

To update your Amazon Q retriever

```
aws qbusiness update-retriever \  
--application-id application-id \  
--retriever-id retriever-id \  
--display-name display-name \  
--role-arn roleArn \  
--configuration kendraIndexConfiguration="{indexId=<kendra-index-id>}"
```

Selecting an Amazon Kendra retriever to an Amazon Q Business application

To select an existing Amazon Kendra retriever to your Amazon Q Business application, you can use the AWS Management Console or the [CreateRetriever](#) API operation.

If you use the API, you select and connect your Amazon Kendra retriever when you use the `CreateRetriever` API operation.

If you use the console, selecting and connecting an Amazon Kendra retriever is a two-step process. This topic provides instructions for the first step: Selecting an Amazon Kendra retriever. For instructions for the second step, see [Connecting an Amazon Kendra retriever to an Amazon Q Business application](#).

Note

If you use an Amazon Kendra retriever, data in your Amazon Kendra will be connected to your Amazon Q application. If you choose this option, you can't use Amazon Q data connectors or direct document upload for your application.

For more information about Amazon Kendra, see the following topics in the Amazon Kendra User Guide and API Reference:

- [What is Amazon Kendra?](#)
- [Creating a data source connector](#)
- [Amazon Kendra API Reference](#)

The following tabs provide a procedure for the AWS Management Console and code samples for the AWS CLI.

Console

To create an Amazon Kendra retriever

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. The, in **Select retriever** choose **Use existing retriever** – Choose an Amazon Kendra index you have previously created as a retriever. All data sources synced to your Amazon Kendra index will be connected to your Amazon Q Business application.
4. In **Tags** – Choose whether you want to add **Retriever tags**.
5. To connect your application to your data sources, choose **Next**.

AWS CLI

To create an Amazon Kendra retriever

```
aws qbusiness create-retriever \  
--display-name display-name \  
--type KENDRA_INDEX \  
--role-arn roleArn \  
--configuration kendraIndexConfiguration="{indexId=<kendra-index-id>
```

Managing Amazon Kendra retrievers

To manage Amazon Kendra retrievers, you can take the following actions:

Actions

- [Deleting an Amazon Kendra retrievers](#)
- [Getting properties of an Amazon Kendra retriever](#)
- [Listing Amazon Kendra retrievers](#)
- [Updating an Amazon Kendra retriever](#)

Deleting an Amazon Kendra retrievers

To delete an Amazon Kendra retriever, you can use the console or the [DeleteRetriever](#) API operation.

If you use the console, the only way to delete your Amazon Kendra retriever from your Amazon Q Business application is to delete your Amazon Q Business application.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To delete an Amazon Kendra retriever

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, choose **Actions**.
3. Choose **Delete**.
4. In the dialog box that opens, type **Delete** to confirm deletion, and then choose **Delete**.

You are returned to the service console while your application is deleted. When the deletion process is complete, the console displays a message confirming successful deletion.

AWS CLI

To delete an Amazon Kendra retriever

```
aws qbusiness delete-retriever \  
--application-id application-id \  
--retriever-id retriever-id
```

Getting properties of an Amazon Kendra retriever

To get the properties of an Amazon Kendra retriever, you can use the console or the [GetRetriever](#) API operation.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To get the properties of an Amazon Kendra retriever

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, select the name of your application from the list of applications.
3. For **Retriever settings**, the following settings are available:
 - **Retriever** – The type of retriever that you're using.
 - **Document count** – The number of documents that are attached to your index.
 - **Last modified time** – The time that your index was last modified.
 - **Index ID** – The ID of the index attached to your retriever.
 - **Storage used** – The amount of storage that your index is using.
 - **Index status** – The status of your index.

Note

You can't edit or update retriever or index settings.

AWS CLI

To get properties of an Amazon Kendra retriever

```
aws qbusiness get-retriever \  
--application-id application-id \  
--retriever-id retriever-id
```

Listing Amazon Kendra retrievers

To list Amazon Kendra retrievers, you can use the console or the [ListRetrievers](#) API operation.

If you use the console, the list of native retrievers and indices attached to them correspond to the list of applications that you have created.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To list Amazon Kendra retrievers

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. For **Applications**, a list of all retrievers (with indices associated) that you have created is available.

AWS CLI

To list Amazon Kendra retrievers

```
aws qbusiness list-retrievers \  
--application-id application-id \  
--max-results maximum-result-to-display
```

Updating an Amazon Kendra retriever

To update your Amazon Kendra retriever, you can use the [UpdateRetriever](#) API operation.

You can't update your Amazon Kendra retriever using the console.

The following tab provides code examples for the AWS CLI.

Console

This action is not supported on the console.

AWS CLI

To update an Amazon Kendra retriever

```
aws qbusiness update-retriever \  
--application-id application-id \  
--retriever-id retriever-id \  
--display-name display-name \  
--role-arn roleArn \  
--configuration kendraIndexConfiguration="{indexId=<kendra-index-d>}"
```

Connecting data sources to an Amazon Q Business application

Important

Starting April 30, 2024, all new applications using [legacy identity management](#) will need to use IAM Identity Center directly to manage user access. All existing Amazon Q applications will need to migrate to using IAM Identity Center for user management by July 31, 2024. We recommend you integrate any new application you're creating directly with IAM Identity Center.

After you select a retriever for your Amazon Q Business application, you connect data sources to it. Available data sources vary based on your choice of the retriever.

If you use an Amazon Q retriever, you can choose from the following options:

- Connect to any Amazon Q supported data source connectors by using the [CreateDataSource](#) API operation.
- Upload documents directly by using the [BatchPutDocument](#) API operation.

If you use an existing Amazon Kendra retriever, only data sources already connected to your Amazon Kendra index are available in your application.

To connect data sources, choose a topic based on your data source preference for your Amazon Q application.

Topics

- [Upload documents](#)
- [Amazon Kendra retriever](#)
- [Amazon Q data source connectors](#)

Upload documents

To upload documents directly to an Amazon Q Business application, you can use the AWS Management Console or the [BatchPutDocument](#) API operation.

If you use an Amazon Kendra index to retrieve your documents, you can't directly upload documents.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To upload documents

Note

This procedure is available if you chose the **Use native retriever** option to configure your application.

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, for **Upload documents**, select one of the following methods to add your files:
 - Drag and drop the document files that you want to upload.
 - Add your documents to the application, and then select **Choose files**.
5. After choosing your files, choose **Upload**.

You are returned to the Amazon Q console while your documents are uploaded. The console displays a confirmation message when your documents are successfully uploaded.

Note

Files can only be uploaded after the Amazon Q retriever and index creation process has completed.

AWS CLI

To upload documents directly

```
aws qbusiness batch-put-document \  
--application-id application-id \  
--index-id index-id \  
--documents documents-to-add \  
--data-source-sync-id data-source-sync-id \  
--role-arn roleArn
```

Delete uploaded documents

To delete documents that have been directly uploaded to an application, you can use the console or the [BatchDeleteDocument](#) API operation. You can delete specific documents or all documents.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To delete specific directly uploaded documents

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, select the name of the application that your uploaded files belong to.
3. From your applications page, from **Data sources**, choose **Uploaded files**.
4. In **Uploaded files**, choose **Document name**, and then select the documents that you want to delete.
5. Choose **Delete files**.

You are returned to the service console while your application is deleted. When the deletion process is complete, the console displays a message confirming successful deletion.

To delete all directly uploaded documents

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.

2. In **Applications**, select the name of the application that your uploaded files belong to.
3. From your applications page, from **Data sources**, select **Uploaded files**.
4. Select **Actions**, and then choose **Delete**.
5. When the deletion process is complete, the console displays a message confirming successful file deletion.

AWS CLI

To delete documents

```
aws qbusiness batch-delete-document \  
--application-id application-id \  
--index-id index-id \  
--documents documents-to-delete \  
--data-source-sync-id data-source-sync-id
```

Connecting an Amazon Kendra retriever to an Amazon Q Business application

To use an Amazon Kendra index as a retriever for Amazon Q Business, you must have already configured an Amazon Kendra index and connected it with data. For more information, see [What is Amazon Kendra?](#) and [Are you a first-time Amazon Kendra user?](#) in the Amazon Kendra Developer Guide.

To add an existing Amazon Kendra retriever to your Amazon Q application, you can use the AWS Management Console or the [CreateRetriever](#) API operation. If you use the console, selecting and connecting an Amazon Kendra retriever is a two-step process. The first step is when you [select an Amazon Kendra retriever](#). In this topic, you perform the second step—connecting an Amazon Kendra retriever.

If you use the API, you create your web experience after connecting your Amazon Kendra retriever using the [CreateWebExperience](#) API operation. If you use the console, connecting your Amazon Kendra retriever also automatically creates your Amazon Q web experience. At the end of the retriever connection process, your Amazon Kendra powered Amazon Q web experience is ready to be previewed, enhanced, and deployed.

Note

If you select an Amazon Kendra retriever, data in your Amazon Kendra is connected to your Amazon Q application.

Console

To connect an Amazon Kendra retriever

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Kendra retriever](#).
4. Then, in **Content sources**, for **Amazon Kendra indexes** – Choose the Amazon Kendra index that you want to use for your Amazon Q Business application. Then, enter the following information:
 - **Service access** – Provide the IAM access role to connect Amazon Kendra to Amazon Q Business. Use an existing role, or create a new one.
 - **Service role name** – Provide a name for your IAM access role. Or, choose to use the auto-generated role that's provided.
5. To connect your Amazon Kendra indexes to the application, choose **Create application**.

You are returned to the Amazon Q console while your web application is created.

AWS CLI

To create and connect an Amazon Kendra retriever

```
aws qbusiness create-retriever \  
--application-id application-id \  
--display-name display-name \  
--type KENDRA_INDEX \  
--role-arn roleArn \  
--configuration kendraIndexConfiguration="{indexId=<kendra-index-id>}"
```

Note

For information on managing your Amazon Kendra retriever, see [Managing Amazon Kendra retrievers](#).

Amazon Q Business data sources

To connect a data source to your Amazon Q Business application, you can use the AWS Management Console or the [CreateDataSource](#) API operation.

By using the `CreateDataSource` API operation, you can configure tags, sync run schedules, and configure Amazon VPC settings. Then, you can use the `configuration` parameter to provide all other configuration information specific to your data source connector.

If you use the console, creating the data source and configuring it are a single step. After your data source is successfully configured and added, Amazon Q automatically creates a Amazon Q web experience for you.

If you use the API, you use the [CreateWebExperience](#) API operation after connecting your data sources to create your web experience.

Note

This procedure is available if you chose the [Use native retriever](#) option to configure your application.

Console

To connect a data source to an Amazon Q application

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).
4. Then, from **Data sources** – Add an available data source to connect your Amazon Q application.

You can add up to 50 data sources.

5. For information on configuring your chosen data source, see [Supported connectors](#) to find configuration information specific to your data source.
6. To connect your configured data source to your application, choose **Add data sources**.

At the end of this step, your Amazon Q web experience is ready to be previewed, enhanced, and deployed.

AWS CLI

To connect a data source

```
aws qbusiness create-data-source \
--application-id application-id \
--index-id index-id \
--configuration data-source-configuration-details \
--display-name display-name \
--role-arn roleArn \
--description description \
--document-enrichment-configuration document-enrichment-configuration \
--sync-schedule sync-schedule-information \
--tags tags \
--vpc-configuration vpc-configuration
```

Managing Amazon Q Business data sources

To manage data source connectors, you can perform the following actions:

Actions

- [Deleting an Amazon Q Business data source connector](#)
- [Getting properties of an Amazon Q Business data source connector](#)
- [Listing Amazon Q Business data source connectors](#)
- [Updating Amazon Q Business data source connectors](#)
- [Starting data source connector sync jobs](#)
- [Stopping data source connector sync jobs](#)
- [Listing data source connector sync jobs](#)

Deleting an Amazon Q Business data source connector

To delete an Amazon Q Business data source connector, you can use the console or the [DeleteDataSource](#) API operation .

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To delete an Amazon Q data source connector

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, select the application that you want to delete data sources from.
3. On the application page, from **Data sources**, select the data source that you want to delete.
4. From **Actions**, choose **Delete**.
5. In the dialog box that opens, type **Delete** to confirm deletion, and then choose **Delete**.

You are returned to the service console while your data source connector is deleted. When the deletion process is complete, the console displays a message confirming successful deletion.

AWS CLI

To delete an Amazon Q data source connector

```
aws qbusiness delete-data-source \  
--application-id application-id \  
--index-id index-id \  
--data-source-id data-source-id
```

Getting properties of an Amazon Q Business data source connector

To get the properties of an Amazon Q Business data source connector, you can use the [GetDataSource](#) API operation.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To get properties of an Amazon Q data source connector

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, select the application you want that contains your data sources.
3. On the application page, from **Data sources**, select the data source that you want to view details for.
4. Under **Data source details**, the following details are available:
 - **Name** – The name of your data source.
 - **Status** – The status of your data source.
 - **Last sync status** – The status of your last sync.
 - **Description** – The description that you gave to your data source.
 - **Type** – The type of data source that you're using.
 - **Last sync time** – The time that your data source was last synced.
 - **Data source ID** – The ID of your data source.
 - **IAM role ARN** – The Amazon Resource Name (ARN) of the IAM role that's associated with your data source.
 - **Current sync state** – The current sync state of your data source.

To get Amazon Q data source connector settings

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, select the application you want that contains your data sources.
3. On the application page, from **Data sources**, select the data source that you want to view details for.
4. For **Data source details**, choose **Settings**.
5. For **Settings**, the following settings are available:

- **IAM role** – The ARN of the IAM that's associated with your data source.
- **Sync scope** – The configuration details for your data source.
- **Sync mode** – The sync type that you chose for your data source.
- **Sync schedule** – The sync schedule that you chose for your data source.
- **Field mappings** – The data source document fields that you chose to map to Amazon Q index fields.

AWS CLI

To get Amazon Q data source connector properties

```
aws qbusiness get-data-source \  
--application-id application-id \  
--index-id index-id \  
--data-source-id data-source-id
```

Listing Amazon Q Business data source connectors

To list Amazon Q Business data source connectors, you can use the console or the [ListDataSources](#) API operation.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To list Amazon Q data source connectors

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, select the application you want that contains your data sources.
3. On the application page, under **Data sources**, a list of data sources connected to your application is displayed.

AWS CLI

To list Amazon Q data source connectors

```
aws qbusiness list-data-sources \  
--application-id application-id \  
--index-id index-id \  
--max-results maximum-number-of-results-to-return
```

Updating Amazon Q Business data source connectors

To update your Amazon Q Business data source connectors, you can use the console or the [UpdateDataSource](#) API operation.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To update a Amazon Q data source connector

Option 1

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, select the application you want to delete data sources from.
3. On the application page, from **Data sources**, select the data source that you want to edit.
4. From **Actions**, choose **Edit**.

You are redirected to your data source configuration page to edit your existing settings.

Option 2

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, select the application you want to delete data sources from.
3. On the application page, from **Data sources**, select the data source that you want to edit.

4. On the data source page, from **Actions**, choose **Edit**.

You are redirected to your data source configuration page to edit your existing settings.

CLI

To update your Amazon Q connector

```
aws qbusiness update-data-source \  
--application-id application-id \  
--data-source-id data-source-id \  
--index-id index-id \  
--configuration data-source-configuration-details \  
--description description \  
--display-name display-name \  
--document-enrichment-configuration document-enrichment-configuration \  
--role-arn roleArn \  
--sync-schedule sync-schedule-information \  
--vpc-configuration vpc-configuration
```

Starting data source connector sync jobs

To start Amazon Q Business data source connector sync jobs, you can use the console or the [StartDataSourceSyncJobs](#) API operation.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To start your Amazon Q data source connector sync jobs

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/aamazonq/>.
2. In **Applications**, select the application you want to sync data sources in.
3. On the application page, from **Data sources**, select the data source that you want to sync.
4. Choose **Sync now**.

The console displays a message confirming that your sync job has started successfully.

Note

You can also view your sync job report in the Amazon CloudWatch console.

AWS CLI

To start your Amazon Q data source connector sync jobs

```
aws qbusiness start-data-source-sync-job \  
--application-id application-id \  
--index-id index-id \  
--data-source-id data-source-id
```

Stopping data source connector sync jobs

To stop your Amazon Q Business connector sync jobs, you can use the console or the [StopDataSourceSyncJobs](#) API operation.

Note

You can only stop a sync job already in progress.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To stop your Amazon Q data source connector sync jobs

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, select the application you want to sync data sources in.
3. On the application page, from **Data sources**, select the data source that you want to stop the sync for.

4. Choose **Stop sync**.
5. In the dialog box that opens, type **Stop** to confirm your action and then select **Stop sync**.

The console displays a message confirming that your data source sync job is being stopped.

AWS CLI

To stop your Amazon Q data source connector sync jobs

```
aws qbusiness stop-data-source-sync-job \  
--application-id application-id \  
--data-source-id data-source-id \  
--index-id index-id
```

Listing data source connector sync jobs

To list Amazon Q Business data source connector sync jobs that are in progress, you can use the console or the [ListDataSourceSyncJobs](#) API operation.

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To list your Amazon Q data source connector sync jobs

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. In **Applications**, select the application you want that contains your data sources.
3. On the application page, from **Data sources**, select the data source that you want to view details for.
4. Under **Data source details**, choose the **Sync run history** tab.

You will see a list of ongoing, completed, and failed sync jobs for your data sources.

CLI

To list your Amazon Q data source connector sync jobs

```
aws qbusiness list-data-source-sync-job \  
--application-id application-id \  
--data-source-id data-source-id \  
--index-id index-id \  
--max-results max-results-to-return
```

Deploying an Amazon Q Business web experience

 Important

Starting April 30, 2024, all new applications using [legacy identity management](#) will need to use IAM Identity Center directly to manage user access. All existing Amazon Q applications will need to migrate to using IAM Identity Center for user management by July 31, 2024. We recommend you integrate any new application you're creating directly with IAM Identity Center.

If you're integrating your Amazon Q application with an external SAML 2.0 compliant identity provider (IdP) (including using IAM Identity Center as your identity provider by creating a [customer managed](#) IAM Identity Center application), you deploy the web experience that you created so your end users can access it. Before you can deploy the web experience, you must set up end user authentication.

For your end users to log in and chat, Amazon Q requires that you integrate your web experience with an identity provider (IdP) that's compliant with SAML 2.0. This integration is required so that only authorized end users from within your organization have access to your content. Amazon Q can work with any IdP that's compliant with SAML 2.0. Amazon Q uses service-initiated single sign-on (SSO) to authenticate users. IdP-initiated SSO is *not* supported.

 Note

During Preview, an Amazon Q application supports only 50 end users. If you need more capacity, contact [Support](#).

To create and deploy your Amazon Q web experience, you can use either the AWS Management Console or the Amazon Q API. If you choose the API, use the [CreateWebExperience](#) API operation to create and deploy your web experience. Then, provide the deployment configuration information using the [WebExperienceAuthConfiguration](#) object.

If you use the console to create your Amazon Q application, a web experience is created automatically. Then, you deploy the web experience by specifying your configuration information on the console. If you use the console, setting up this connection involves copying and entering information from the Amazon Q console into the IdP console, and the other way around.

Topics

- [Overview of integrating Amazon Q Business with an Identity Provider \(IdP\)](#)
- [Key IdP integration concepts](#)
- [Steps for deploying your Amazon Q Business web experience](#)
- [Troubleshooting Amazon Q Business and identity provider integration](#)

Overview of integrating Amazon Q Business with an Identity Provider (IdP)

To deploy an Amazon Q Business web experience, you must set up end user authentication. Amazon Q requires that you integrate your web experience with an identity provider (IdP) that's compliant with SAML 2.0. This integration is required so that only authorized end users from within your organization have access to your content. For more information, see [Deploying an Amazon Q Business web experience](#).

The following gives you a high-level overview of the required steps to integrate Amazon Q Business with your IdP:

1. Create your Amazon Q web experience.
2. Create a new app integration in your IdP.
3. Share your Amazon Q configuration information with your IdP. This step starts the IdP and Amazon Q connection configuration process.
4. Share your IdP's federation metadata with Amazon Q. This step establishes a trust relationship between your IdP and Amazon Q. The trust relationship allows Amazon Q to validate user information that's communicated by your IdP. Establishing this trust relationship ensures that only a user who has permissions to access your application can access it.

5. Share the email attribute name (required) and group attribute name (optional) from your IdP with Amazon Q. Amazon Q uses this information to perform document access control based on the user's identity. This step ensures that your authenticated end user only sees chat responses generated from documents they have access to.

For more information about the terms used in describing the integration process, see [Key IdP integration concepts](#).

Topics

- [Overview of deploying Amazon Q web experience steps](#)

Overview of deploying Amazon Q web experience steps

This section gives you an overview of the process of deploying your web experience by using the AWS Management Console and the AWS CLI.

As a prerequisite, make sure you completed creating your application.

For more in-depth, customized instructions to guide you through deploying your web experience using specific IdPs, choose from the following options:

- [Using IAM Identity Center](#)
- [Using Entra ID](#)
- [Using Okta](#)
- [Using PingIdentity](#)

The following tabs provide a procedure for the AWS Management Console and code examples for the AWS CLI.

Console

To deploy your Amazon Q web experience

1. Sign in to the AWS Management Console and open the Amazon Q console at <https://console.aws.amazon.com/amazonq/>.
2. Complete the steps to [create your Amazon Q application](#).
3. Complete the steps for [selecting an Amazon Q retriever](#).

4. Complete the steps for [connecting data sources](#).
5. Optionally, complete the steps for [enhancing an application](#).
6. Optionally, complete the steps to [customize your web experience](#).
7. Then, in **Applications**, select your application, and choose **Deploy web experience**.
8. In **Service access**, enter the following information:
 - **Service access** – A service access role assumed by end users when they sign in to your web experience that grants them permission to start and manage conversations Amazon Q. You can choose to use an existing role or create a new role.
 - **Service role name** – A name for the service role you created for easy identification on the console.
9. From [Identity provider](#), copy the following information to provide to the IdP you're using:
 - [Assertion consumer service \(ACS\) URL](#) – Copy the ACS URL and enter it in the relevant section of your IdP.
 - [Audience URI \(SP Entity ID\)](#) – Copy the Audience URI (SP Entity ID) and enter it in the relevant section of your IdP.
10. In **Provide metadata from your IdP**, enter the following information:
 - Upload the [metadata generated by your IdP as an XML file](#) using **Import from XML**.

See [Key IdP integration concepts](#) and [Integration process overview](#) for more details.
11. In **Configure user and group mapping**, enter the following information to allow ACLs to be active for end users using the web experience:
 - [Email attribute of SAML assertion](#) – Provide the attribute name that maps to user email.
 - [User group field attribute of SAML assertion - optional](#) – Provide the attribute name that maps to user groups.

See [Key IdP integration concepts](#) and [Integration process overview](#) for more details.
12. To finish deploying your web experience, choose **Deploy**.

You are redirected to the Amazon Q control panel while your web experience deployment process finishes. After your application is deployed, your end users can access and chat in the web experience using the deployed web experience URL that's generated in the web experience details page by Amazon Q.

AWS CLI

To deploy a web experience

```
aws qbusiness create-web-experience \  
--application-id application-id \  
--metadata-xml metadata-xml \  
--role-arn roleArn \  
--user-id-attribute user-id-attribute \  
--user-group-attribute user-group-attribute
```

Key IdP integration concepts

Amazon Q Business requires that you integrate your web experience with an identity provider (IdP) that's compliant with SAML 2.0. This integration is required so that only authorized end users from within your organization have access to your content. For more information, see [Deploying an Amazon Q Business web experience](#). The following are key concepts that will help you understand the terms you encounter during the integration process.

Topics

- [Authorization](#)
- [Authentication](#)
- [Identity provider \(IdP\)](#)
- [Service provider \(SP\)](#)
- [Security Assertion Markup Language \(SAML\)](#)
- [Service provider-initiated single sign-on \(SSO\) flow](#)
- [Identity provider-initiated single sign-on \(SSO\) flow](#)
- [Assertion consumer service \(ACS\) URL](#)
- [Audience URI \(SP entity ID\)](#)
- [XML metadata file](#)
- [SAML assertion](#)
- [Email attribute of SAML assertion](#)
- [User group attribute of SAML assertion](#)

Authorization

Authorization allows a user permissions to access specific resources.

Authentication

Authentication confirms a user's identity—that users are who they say they are.

Identity provider (IdP)

An identity provider (IdP) is a service that stores, manages, maintains, and verifies user identities for your application (in this case, Amazon Q). Some examples of IdPs are AWS IAM Identity Center, Okta, and Microsoft EntraID.

Service provider (SP)

A service provider (SP) is any entity—in this case, Amazon Q—that requests user authentication and authorization services from an IdP. Amazon Q takes the authentication information received from an IdP and uses it to authorize the end user's web experience session based on user authorization levels.

Security Assertion Markup Language (SAML)

SAML is an XML-based standard for transferring user identity data between the service provider (SP)—in this case, Amazon Q—and an identity provider (IdP) such as Okta, Ping, or EntraID. SAML supports two types of sign-in flows: Service initiated and IdP initiated.

Amazon Q only supports IdPs that are compliant with SAML 2.0.

Service provider-initiated single sign-on (SSO) flow

A SAML flow in which a service provider (SP) initiates the sign-in process.

Important

Amazon Q uses service-initiated single sign-on (SSO) to authenticate users. IdP-initiated SSO is *not* supported.

Identity provider-initiated single sign-on (SSO) flow

A SAML flow in which the identity provider (IdP) (for example, Okta) initiates the sign-in process.

⚠ Important

Amazon Q doesn't support IdP-initiated SSO.

Assertion consumer service (ACS) URL

An assertion consumer service (ACS) URL is an endpoint on the service provider (SP)—in this case, Amazon Q—where the IdP redirects its authentication response. This endpoint decides where your IdP sends its SAML response after authenticating a user.

Audience URI (SP entity ID)

The audience URI (service provider entity ID) is the unique ID of your service provider (SP). An identity provider (IdP) uses the audience URI to identify and direct its SAML response to a service provider.

XML metadata file

The XML metadata file is the document that contains the configuration information generated by your IdP during your SP-initiated single sign-on (SSO) process. The document contains the information needed for your SP and your IdP to trust and communicate with each other.

SAML assertion

A SAML assertion is a message that's exchanged between your SP and your IdP that confidentially identifies a user. Assertions contain information about user identity, their group membership, the information that users can access, and any other relevant information.

Email attribute of SAML assertion

The email attribute of a SAML assertion is the attribute that your IdP maps user email to. For example, a user email address of *mary_major@example.com* can be mapped to the attribute `user_email`. Amazon Q uses this attribute value to resolve user access level to documents.

User group attribute of SAML assertion

The user group attribute of a SAML assertion is the attribute that the IdP maps user groups to. For example, the user groups "Research" and "Science" can be mapped to the attribute `user_group`. Amazon Q uses this attribute value to resolve user access level to documents.

Steps for deploying your Amazon Q Business web experience

To deploy your Amazon Q Business web experience to your end users, you must integrate your Amazon Q application with an identity provider (IdP) that's compliant with SAML 2.0. You do this during the [deploy your web experience](#) process.

To integrate your external SAML 2.0-compliant IdP, you must switch between tasks on the Amazon Q console and your IdP account.

This section guides you through the process of deploying your web experience using the following IdPs. You can use similar steps for integrating your Amazon Q application with any IdP that's compliant with SAML 2.0.

- [Using IAM Identity Center](#)
- [Using Entra ID](#)
- [Using Okta](#)
- [Using PingIdentity](#)

Note

As a prerequisite, make sure you've completed [creating your application](#).

Important

Amazon Q uses service-initiated single sign-on (SSO) to authenticate users. IdP-initiated SSO is *not* supported.

Topics

- [Setting up Amazon Q Business with IAM Identity Center as identity provider](#)
- [Setting up Amazon Q Business with Microsoft Entra ID as identity provider](#)
- [Setting up Amazon Q Business with Okta as identity provider](#)
- [Setting up Amazon Q Business with PingIdentity as identity provider](#)

Setting up Amazon Q Business with IAM Identity Center as identity provider

The following steps show how to set up Amazon Q Business with AWS IAM Identity Center as your SAML 2.0-compliant identity provider (IdP). Integrating Amazon Q with IAM Identity Center requires that you switch between tasks on the Amazon Q console and the IAM Identity Center console.

Prerequisites

Before you start to integrate Amazon Q with IAM Identity Center, make sure that you have completed the following tasks:

- Created an Amazon Q Business application, selected a retriever, added your desired data sources, and previewed an Amazon Q Business web experience.
- Enabled an IAM Identity Center instance, provisioned at least one user, and provided each user with a valid email address. For more details, see [Configure user access with the default IAM Identity Center directory](#).

Note

To deploy your web experience using IAM Identity Center as an identity provider, Amazon Q requires you to create a custom application. This is because Amazon Q is not an AWS managed application. However, IAM Identity Center account instances can't support custom IAM Identity Center applications. So, you need to use an IAM Identity Center organizations instance (which supports custom applications) to integrate IAM Identity Center with Amazon Q. For more information on IAM Identity Center instances, see [IAM Identity Center capabilities](#).

To integrate Amazon Q with IAM Identity Center

1. In the Amazon Q console, choose the Amazon Q application you want to integrate with IAM Identity Center.
2. On the **Applications** page, from **Applications**, choose the application you want to deploy. Then, choose **Deploy web experience**.

The screenshot shows the Amazon Q Applications console. At the top, there's a navigation bar with 'Applications' and a red callout '1' pointing to the 'Applications' tab. Below that, there's a 'Documentation' link. The main content area is titled 'Applications' and includes a 'How it works' section with four steps: 'Create generative AI application', 'Enhance application - optional', 'Customize web experience', and 'Deploy web experience'. The 'Deploy web experience' button is highlighted with a red box and a red callout '3'. Below this, there's a table of applications with columns for Name, Application status, Retriever, Creation time, Web experience status, and Deployed URL. The first row, 'GenAI-application-tk964', is highlighted with a red box and a red callout '2'.

3. On the **Deploy web experience** page, for **Service access**, choose to **Create a use a new service role** or **Use an existing service role**. If you choose to create a new service role, Amazon Q, will automatically create a name for it.

The screenshot shows the 'Service access' configuration page. The title is 'Service access' with an 'Info' link. Below the title, it says 'Amazon Q requires permissions to use other services on your behalf.' The main section is 'Choose a method to authorize Amazon Q', which is highlighted with a red box. It contains two radio button options: 'Create and use a new service role' (selected) and 'Use an existing service role'. Below this is a 'Service role name' input field.

4. In the **Configure your [Identity provider](#)** section, do the following:
 - Copy the **[Assertion consumer service\(ACS\) URL](#)** displayed on the console to a text editor of your choice
 - Copy the **[Audience URI \(SP EntityID\)](#)** displayed on the console to a text editor of your choice.

Configure your identity provider (IdP) [Info](#)

Configure your SAML support on your IdP. Provide the Assertion Consumer Service (ACS) URL and Audience URI information below to your IdP. This causes IdP metadata to be generated.

Application consumer service (ACS) URL

The ACS URL is the endpoint where the SAML response will be sent.

Audience URI (SP Entity ID)

Determines the intended recipient or audience for the SAML assertion.

You will use this information later in this procedure.

- Then, switch to the [IAM Identity Center console](#).
- From the IAM Identity Center console, from the left navigation pane, expand **Application assignments** and choose **Applications**.

The screenshot displays the IAM Identity Center Dashboard. On the left, a navigation sidebar is visible with the following items: 'IAM Identity Center' (with a close icon), 'Managing instance AmazonQ', 'Dashboard' (highlighted in blue), 'Users', 'Groups', 'Settings', 'Multi-account permissions' (with a dropdown arrow), 'AWS accounts', 'Permission sets', 'Application assignments' (with a dropdown arrow), and 'Applications' (highlighted with a red box). Below these are 'Related consoles' including 'AWS Organizations' and 'IAM'. The main content area is titled 'IAM Identity Center > Dashboard' and 'Dashboard'. It contains three main sections: 1. 'IAM Identity Center setup' with a 'Confirm your identity source' section (including a 'Confirm identity source' button), a 'Manage permissions for multiple AWS accounts' section (including a 'Manage permissions' button), and a 'Set up application user and group assignments' section (including a 'Set up applications' button). 2. 'Settings summary' with a 'Go to settings' button and fields for 'Instance name - AmazonQ', 'Identity source - Identity Center directory', 'Region', 'Organization ID', 'AWS access portal URL - Edit', and 'Issuer URL'. 3. 'What's new' with a link for 'Introducing trusted identity'.

7. On the **Applications** page, from **Customer managed**, choose **Add application**.

Applications

Administer users and groups for AWS managed or customer managed applications that support identity federation with SAML 2.0 or OAuth 2.0.

[Learn more](#)

Add application

AWS managed | **Customer managed**

Customer managed applications (2) Actions ▾

A customer managed application can be selected from the IAM Identity Center application catalog or manually configured.

Search for a customer managed application

Show incomplete configurations < 1 > ⚙

Application ▾	Owning account ID ▾	Date created ▾	Certificate expiration	Status ▾

- On the **Select application type** page, for the **Setup preference**, choose **I have an application I want to set up**.

Select application type

You can set up an application that you already have to work with IAM Identity Center, or you can select an application from the IAM Identity Center application catalog. Applications in the catalog are already set up to work with IAM Identity Center.

-  To set up an AWS managed application to work with IAM Identity Center, you must configure the application directly from the console for the applicable service. [View all AWS services](#) 

Setup preference

If you already have an application, you can set it up to use OAuth 2.0 and OIDC for trusted identity propagation, or SAML 2.0 for identity federation. Applications in the IAM Identity Center application catalog support SAML 2.0 only.

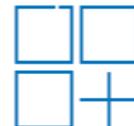
I have an application I want to set up

Manually set up your application to work with IAM Identity Center. You can configure your application to use OAuth 2.0 and OIDC for trusted identity propagation or SAML 2.0 for identity federation. [Learn more about manually setting up applications](#) 



I want to select an application from the catalog

Select an application from a catalog of commonly used applications that are already set up to work with IAM Identity Center. These applications support SAML 2.0 for identity federation. [Learn more about the IAM Identity Center application catalog](#) 



9. In the **Application type** section, choose **SAML 2.0** and choose **Next**.

Application type

OAuth 2.0

This application can be set up for trusted identity propagation. The application uses OAuth 2.0 token exchange to authorize access to other trusted applications on behalf of its users. OIDC is used to authenticate users.

SAML 2.0

This application supports SAML 2.0 for identity federation only. Trusted identity propagation isn't supported. The application exchanges data in XML SAML format to authenticate users and authorize access to resources.

1

Available capabilities	OAuth 2.0	SAML 2.0
Single sign-on	✓	✓
Assign users and groups	✓	✓
Trusted identity propagation What is this?	✓	-
OpenID Connect (OIDC) What is this?	✓	-

2

Next

10. On the **Configure application** page, in **Display name** enter a name for your application. Optionally, enter a description in **Description**.

Configure application

Display name

Description

The description you type here does not appear in the AWS access portal. However, it will be visible in the IAM Identity Center console and when using IAM Identity Center APIs.

11. In the **IAM Identity Center metadata** section, choose **Download** to download the IAM Identity Center [SAML metadata](#) file. You will need this when you return to the Amazon Q console.

IAM Identity Center metadata

Your cloud application may require the following certificate and metadata details to recognize IAM Identity Center as the identity provider.

IAM Identity Center SAML metadata file

 [Download](#)

 [https://portal.sso.us-east-](https://portal.sso.us-east-1.amazonaws.com/saml/metadata/OTcyNjM1NDI5Nzc5X2lucy0yNGZmNzBkMDczNzE1ODgy)

[1.amazonaws.com/saml/metadata/OTcyNjM1NDI5Nzc5X2lucy0yNGZmNzBkMDczNzE1ODgy](https://portal.sso.us-east-1.amazonaws.com/saml/metadata/OTcyNjM1NDI5Nzc5X2lucy0yNGZmNzBkMDczNzE1ODgy)

IAM Identity Center sign-in URL

 [https://portal.sso.us-east-](https://portal.sso.us-east-1.amazonaws.com/saml/assertion/OTcyNjM1NDI5Nzc5X2lucy0yNGZmNzBkMDczNzE1ODgy)

[1.amazonaws.com/saml/assertion/OTcyNjM1NDI5Nzc5X2lucy0yNGZmNzBkMDczNzE1ODgy](https://portal.sso.us-east-1.amazonaws.com/saml/assertion/OTcyNjM1NDI5Nzc5X2lucy0yNGZmNzBkMDczNzE1ODgy)

IAM Identity Center sign-out URL

 <https://portal.sso.us-east-1.amazonaws.com/saml/logout/OTcyNjM1NDI5Nzc5X2lucy0yNGZmNzBkMDczNzE1ODgy>

IAM Identity Center SAML issuer URL

 [https://portal.sso.us-east-](https://portal.sso.us-east-1.amazonaws.com/saml/assertion/OTcyNjM1NDI5Nzc5X2lucy0yNGZmNzBkMDczNzE1ODgy)

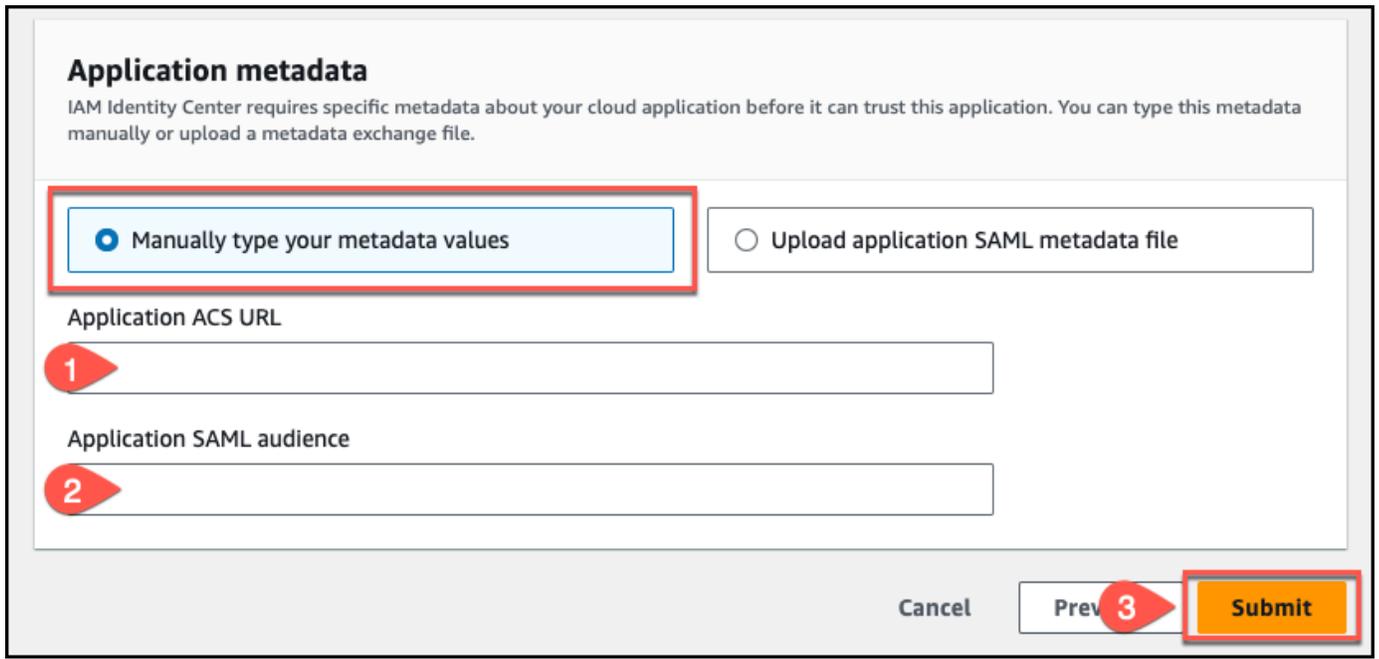
[1.amazonaws.com/saml/assertion/OTcyNjM1NDI5Nzc5X2lucy0yNGZmNzBkMDczNzE1ODgy](https://portal.sso.us-east-1.amazonaws.com/saml/assertion/OTcyNjM1NDI5Nzc5X2lucy0yNGZmNzBkMDczNzE1ODgy)

IAM Identity Center Certificate

 [Download](#)

12. In **Application properties**, (if you're configuring access to the Amazon Q application directly from your IdP's application portal instead of using the deployed Amazon Q web experience), you must choose to specify the deployed experience URL from the Amazon Q console as the **Application start URL**.
13. Scroll down to the **Application metadata** section, and choose **Manually type your metadata values**.
14. Then, do the following:
 - For **Application (ACS) URL** – Enter the [Assertion consumer service\(ACS\) URL](#) value you copied from the Amazon Q console.
 - **Application SAML audience URI** – Enter the [Audience URI \(SP EntityID\)](#) value you copied from the Amazon Q console.

Then, choose **Submit**.



Application metadata

IAM Identity Center requires specific metadata about your cloud application before it can trust this application. You can type this metadata manually or upload a metadata exchange file.

Manually type your metadata values Upload application SAML metadata file

Application ACS URL

1

Application SAML audience

2

Cancel Prev 3 Submit

15. On the **Custom SAML 2.0 application** application page, scroll down to the **Assigned users and groups** section and choose **Assign users and groups**.

Custom SAML 2.0 application

Details



Display name
Custom SAML 2.0 application

Service
Custom SAML 2.0 application

Description
Custom SAML 2.0 application

Status
Active

Application ARN
[Redacted]

Owning account ID
-

User and group assignments
Require assignments

Authentication with trusted token issuer
 Not configured

Date created
February 26, 2024

Assigned users and groups (1/1)

The following users and groups can access this application. [Learn more about user and group assignments](#)

< 1 > ⚙

<input checked="" type="checkbox"/>	Username of user or group	Type
<input checked="" type="checkbox"/>	[Redacted]	User

16. On the **Assign users to Custom SAML 2.0 application** table, select one or more users for your application and then choose **Assign users** to finish assigning users.

Assign users to Custom SAML 2.0 application

 Users you assign here must also have equivalent accounts in the Custom SAML 2.0 application before they can have multi-account access to the application from the AWS access portal. You can create these accounts manually or enable just-in-time (JIT) provisioning in the application to create these accounts automatically.

You can search for the users and groups to grant multi-account access. You can select more than one user or group. [Learn more](#) 

Users (1) | Groups (0)

Users (1)

< 1 > 

 Name | Email

▶ Selected users and groups (1)

Cancel

 **Assign users**

17. From the **Details** pane, choose **Actions** and then choose **Edit attribute mappings**.

 Configuration for 'Custom SAML 2.0 application' has been saved. You must configure attribute mappings for IAM Identity Center to work. 

[IAM Identity Center](#) > [Applications](#) > Custom SAML 2.0 application

Custom SAML 2.0 application

Details

 **Display name**
Custom SAML 2.0 application

Service
Custom SAML 2.0 application

Description
Custom SAML 2.0 application

Status
 Inactive

Application ARN


Owning account ID
-

User and group assignments
Require assignments

Authentication with trusted token issuer
 Not configured

Date created


Actions ▲

- Edit configuration
- Edit attribute mappings**

18. On the **Attribute mappings for Custom SAML 2.0 application** page, do the following:

- Leave the **User attribute in the application** column set to the default attribute name **Subject**.
- For **Maps to this string value or user attribute in IAM Identity Center** – Map the Subject to the email attribute, for example, `${user:email}`. Make sure that the attribute you provide is included in [Supported IAM Identity Center attributes](#).
- Set the **Format** to **unspecified**.

IAM Identity Center > Applications > Custom SAML 2.0 application > Attribute mappings

Attribute mappings for Custom SAML 2.0 application

Attributes you map here become part of the SAML assertion that is sent to the application. You can choose which user attributes in your application map to corresponding user attributes in your connected directory. [Learn more](#)

User attribute in the application	Maps to this string value or user attribute in IAM Identity Center	Format
Subject	<code>\${user:email}</code>	unspecified

[Add new attribute mapping](#)

[Cancel](#) [Save changes](#)

- Choose **Add new attribute mapping**.

19. Then, on the **Attribute mappings for Custom SAML 2.0 application** page, add another attribute mapping by completing the following steps:

- For **User attribute in the application**, enter a name for the attribute, for example, `Email`. Make a note of this attribute name for use later.
- For **Maps to this string value or user attribute in IAM Identity Center** – Enter an attribute or a value that you want to map to the attribute name.

For example, you might want to map the attribute name `Email` with the users email attribute `${user:email}`.

Make sure that the attribute you provide is included in [Supported IAM Identity Center attributes](#).

- c. Set the **Format** to **unspecified**.
- d. Choose **Save changes**.

IAM Identity Center > Applications > Custom SAML 2.0 application > Attribute mappings

Attribute mappings for Custom SAML 2.0 application

Attributes you map here become part of the SAML assertion that is sent to the application. You can choose which user attributes in your application map to corresponding user attributes in your connected directory. [Learn more](#)

User attribute in the application	Maps to this string value or user attribute in IAM Identity Center	Format
Subject	<code>\${user:email}</code>	unspecified
Email	<code>\${user:email}</code>	unspecified

[Add new attribute mapping](#)

[Cancel](#) [Save changes](#)

20. Go back to the Amazon Q console, and make sure you're on the **Deploy web experience** page.
21. Scroll down to the **Provide metadata from your IdP** section. To upload the metadata XML file that you saved in your previous steps, choose **Import from XML**.

Provide metadata from your IdP [Info](#)

Upload your IdP metadata file from an XML file using the button below.

[Import from XML file](#)

File needs to be a valid UTF-8 XML document.

22. In the **Configure user and group mapping** section, do the following:

- For [Email attribute of SAML assertion](#) – Enter the attribute name that you provided in the IAM Identity Center console. For example, **Email** could be an attribute name.

Configure user and group mapping [Info](#)

Provide the fields of the SAML assertion from your IdP so that ACLs can be active for end users using the web experience.

Email attribute of SAML assertion
Provide the attribute name that maps to user email

Email

User group attribute of SAML assertion - optional
Provide the attribute name that maps to user groups.

Enter text

Cancel **Deploy**

Note

Make sure there are no spaces at the end of **Email**.

- For [User group field attribute of SAML assertion - optional](#) – Enter an optional user group attribute.
23. Choose **Deploy**.
 24. Once deployment finishes, a URL should appear on your Amazon Q application page under **Deployed URL**.
 25. Choose the URL to open your Amazon Q web experience and enter credentials for a user that has access to the web experience.

If you encounter HTTP status code 403 (Forbidden) errors, see [Troubleshooting Amazon Q Business and identity provider integration](#).

Setting up Amazon Q Business with Microsoft Entra ID as identity provider

The following steps show how to set up Amazon Q Business with Microsoft Entra ID (formerly Azure Active Directory) as your SAML 2.0-compliant identity provider. Integrating Amazon Q with Entra ID requires that you switch between tasks on the Amazon Q console and in the Entra ID portal.

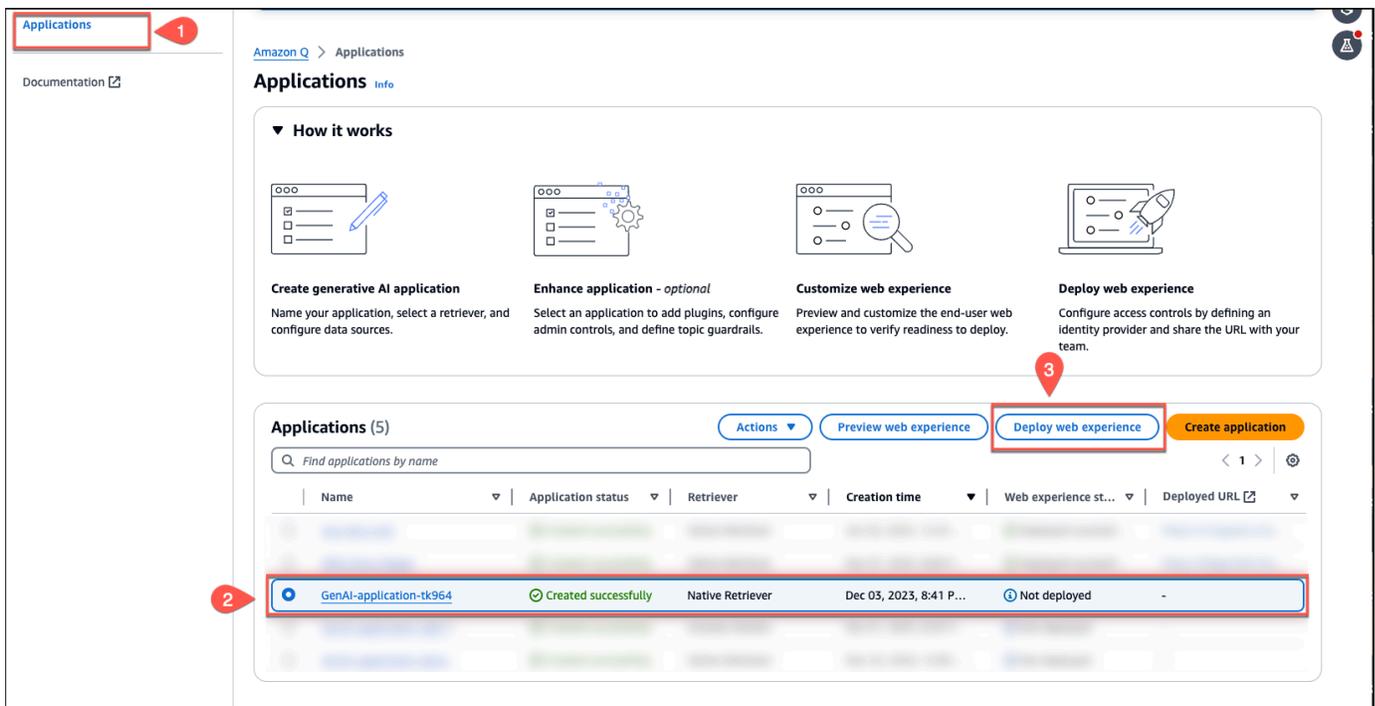
Prerequisites

Before you start to integrate Amazon Q with Entra ID, make sure that you have completed the following tasks:

- Created an Amazon Q Business application, selected a retriever, added your desired data sources, and previewed Amazon Q Business web experience.
- Created an Entra ID instance, provisioned at least one user, and provided each user with a valid email address.

To integrate Amazon Q with Entra ID

1. In the Amazon Q console, choose the Amazon Q application you want to integrate with Entra ID.
2. On the **Applications** page, from **Applications**, choose the application you want to deploy. Then, choose **Deploy web experience**.



3. On the **Deploy web experience** page, for **Service access**, choose to **Create a use a new service role** or **Use an existing service role**. If you choose to create a new service role, Amazon Q, will automatically create a name for it.

Service access [Info](#)

Amazon Q requires permissions to use other services on your behalf.

Choose a method to authorize Amazon Q

Create and use a new service role

Use an existing service role

Service role name

4. In the **Configure your [Identity provider](#)** section, do the following:
 - Copy the **[Assertion consumer service\(ACS\) URL](#)** displayed on the console to a text editor of your choice
 - Copy the **[Audience URI \(SP EntityID\)](#)** displayed on the console to a text editor of your choice.

Configure your identity provider (IdP) [Info](#)

Configure your SAML support on your IdP. Provide the Assertion Consumer Service (ACS) URL and Audience URI information below to your IdP. This causes IdP metadata to be generated.

Application consumer service (ACS) URL

The ACS URL is the endpoint where the SAML response will be sent.

 1

Audience URI (SP Entity ID)

Determines the intended recipient or audience for the SAML assertion.

 2

You will use this information later in this procedure.

5. Then, switch to the Entra ID portal. In the left navigation pane, choose **Enterprise applications**, and then choose **Add**.

The screenshot shows the Microsoft Azure portal interface for Microsoft Entra ID. The top navigation bar includes the Microsoft Azure logo and a search bar. The left-hand navigation pane lists various management options, with 'Enterprise applications' highlighted in a red box. The main content area displays the 'Overview' page for a tenant, featuring a search bar and a table of basic information. A red arrow points from the 'Enterprise applications' link in the navigation pane to the '+ Add' button at the top of the main content area.

6. On the **All applications** page, choose **New application**.

The screenshot shows the Microsoft Azure portal interface for the 'Enterprise applications | All applications' page. The top navigation bar includes the Microsoft Azure logo and a search bar. The left-hand navigation pane lists various management options, with 'All applications' highlighted. The main content area displays the 'Overview' page for all applications, featuring a search bar and a table of basic information. The 'New application' button is highlighted with a red box.

7. In the **Browse Microsoft Entra Gallery** page, choose **Create your own application**.

Microsoft Azure

Search resources, services, and docs (G+)

Home > Default Directory | Enterprise applications > Enterprise applications | All applications >

Browse Microsoft Entra Gallery

+ Create your own application Got feedback?

The Microsoft Entra App Gallery is a catalog of thousands of apps that make it easy to deploy and configure single sign-on (SSO) and automated user provisioning. When deployed, prebuilt templates to connect your users more securely to their apps. Browse or create your own application here. If you are wanting to publish an application you have developed for your organization to discover and use, you can file a request using the process described in [this article](#).

Search application

Single Sign-on : All User Account Management : All Categories : All

Cloud platforms

- Amazon Web Services (AWS)
- Google Cloud Platform
- Oracle

8. Enter a name for your application, choose **Integrate any other application you don't find in the gallery (Non gallery)**, and choose **Create**. It might take a few minutes for your application to be provisioned.

Create your own application



 Got feedback?

If you are developing your own application, using Application Proxy, or want to integrate an application that is not in the gallery, you can create your own application here.

What's the name of your app?

Input name 

What are you looking to do with your application?

- Configure Application Proxy for secure remote access to an on-premises application
- Register an application to integrate with Microsoft Entra ID (App you're developing)
- Integrate any other application you don't find in the gallery (Non-gallery)

Create

9. On the **Application overview** page, in the **Getting started** section, choose **Set up single sign on**.

Getting Started

 **1. Assign users and groups**
Provide specific users and groups access to the applications
[Assign users and groups](#)

 **2. Set up single sign on**
Enable users to sign into their application using their Microsoft Entra credentials
[Get started](#)

 **3. Provision User Accounts**
Automatically create and delete user accounts in the application
[Get started](#)

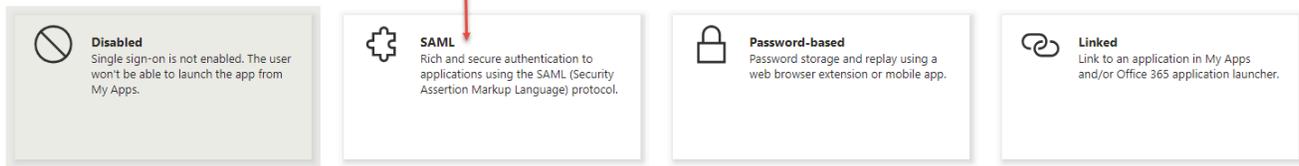
 **4. Conditional Access**
Secure access to this application with a customizable access policy.
[Create a policy](#)

 **5. Self service**
Enable users to request access to the application using their Microsoft Entra credentials
[Get started](#)

10. In the **Select a single sign-on method** pane, choose **SAML**.

Single sign-on (SSO) adds security and convenience when users sign on to applications in Microsoft Entra ID by enabling a user in your organization to sign in to every application they use with only one account. Once the user logs into an application, that credential is used for all the other applications they need access to. [Learn more](#).

Select a single sign-on method [Help me decide](#)



11. In the **Basic SAML Configuration** section, choose **More** (three dots) and then choose **Edit**.

Set up Single Sign-On with SAML

An SSO implementation based on federation protocols improves security, reliability, and end user experiences and is easier to implement. Choose SAML single sign-on whenever possible for existing applications that do not use OpenID Connect or OAuth. [Learn more](#).

Read the [configuration guide](#) for help integrating

1 Basic SAML Configuration Edit

Identifier (Entity ID)	Required
Reply URL (Assertion Consumer Service URL)	Required
Sign on URL	<i>Optional</i>
Relay State (Optional)	<i>Optional</i>
Logout Url (Optional)	<i>Optional</i>

12. Choose **Add identifier**. Then enter the following information:

- For the **Identifier (Entity ID)** field, enter the **Audience URI (SP Entity ID)** that you copied from the Amazon Q console.
- Next, choose **Add reply URL**.
- For the **Reply URL (Assertion Consumer Service URL)** field, enter the **Application consumer service (ACS) URL** that you copied from the Amazon Q console.
- Leave the rest of the fields blank. Choose **Save**.

Basic SAML Configuration

 Save |  Got feedback?

Identifier (Entity ID) * ⓘ

The unique ID that identifies your application to Microsoft Entra ID. This value must be unique across all applications in your Microsoft Entra tenant. The default identifier will be the audience of the SAML response for IDP-initiated SSO.

Default



[Add identifier](#)

Reply URL (Assertion Consumer Service URL) * ⓘ

The reply URL is where the application expects to receive the authentication token. This is also referred to as the "Assertion Consumer Service" (ACS) in SAML.

Index

Default



[Add reply URL](#)

13. On the **Set up single sign-on with SAML** page, scroll down to the **SAML Certificates** section. Download the **Federation Metadata XML** file and save it in your local drive.

1 Basic SAML Configuration  Edit

Identifier (Entity ID)	Required
Reply URL (Assertion Consumer Service URL)	Required
Sign on URL	<i>Optional</i>
Relay State (Optional)	<i>Optional</i>
Logout Url (Optional)	<i>Optional</i>

2 Attributes & Claims

givenname	user.givenname
surname	user.surname
emailaddress	user.mail
name	user.userprincipalname
Unique User Identifier	user.userprincipalname

3 **SAML Certificates**

Token signing certificate  Edit

Status	Active
Thumbprint	
Expiration	8/28/2028, 1:02:40 PM
Notification Email	
App Federation Metadata Url	https://login.microsoftonline.com/888d0b57-69f1... 
Certificate (Base64)	Download
Certificate (Raw)	Download
Federation Metadata XML	Download

Verification certificates (optional)  Edit

Required	No
Active	0
Expired	0

14. In the **Attributes & Claims** section, choose **More** (three dots) and then choose **Edit**.

Set up Single Sign-On with SAML

An SSO implementation based on federation protocols improves security, reliability, and end user experiences and is easier to implement. Choose SAML single sign-on whenever possible for existing applications that do not use OpenID Connect or OAuth. [Learn more.](#)

Read the [configuration guide](#) for help integrating

1

Basic SAML Configuration

Identifier (Entity ID)

Reply URL (Assertion Consumer Service URL) acs

Sign on URL *Optional*

Relay State (Optional) *Optional*

Logout Url (Optional) *Optional*

Edit

2

Attributes & Claims

givenname	user.givenname
surname	user.surname
emailaddress	user.mail
name	user.userprincipalname
Unique User Identifier	user.userprincipalname

Edit

15. In the **Attributes & Claims** page, choose **Unique User Identifier (Name ID)**.

Attributes & Claims ...

+ Add new claim + Add a group claim ☰ Columns | 🗨️ Got feedback?

Required claim

Claim name	Type	Value
Unique User Identifier (Name ID)	SAML	user.userprincipalname [...]

Additional claims

Claim name	Type	Value
http://schemas.xmlsoap.org/ws/2005/05/identity/claims/emailadd...	SAML	user.mail ...
http://schemas.xmlsoap.org/ws/2005/05/identity/claims/givenname	SAML	user.givenname ...
http://schemas.xmlsoap.org/ws/2005/05/identity/claims/name	SAML	user.userprincipalname ...
http://schemas.xmlsoap.org/ws/2005/05/identity/claims/surname	SAML	user.surname ...

∨ Advanced settings

16. In the **Manage claim** page, expand **Choose name identifier format**. For the **Name identifier format** field, select **Unspecified**. Choose **Save**.

Manage claim ...

📄 Save ✕ Discard changes | 🗨️ Got feedback?

Name

Namespace

∧ Choose name identifier format

Name identifier format *

Source * Attribute Transformation Directory schema extension

Source attribute *

∨ Claim conditions

∨ Advanced SAML claims options

17. In the **Attributes & Claims** page, choose **Add new claim**.

Attributes & Claims ...

+ Add new claim
+ Add a group claim
☰ Columns
|
🗨️ Got feedback?

Required claim

Claim name	Type	Value
Unique User Identifier (Name ID)	SAML	user.userprincipalname [...]

Additional claims

Claim name	Type	Value
http://schemas.xmlsoap.org/ws/2005/05/identity/claims/emailadd...	SAML	user.mail ...
http://schemas.xmlsoap.org/ws/2005/05/identity/claims/givenname	SAML	user.givenname ...
http://schemas.xmlsoap.org/ws/2005/05/identity/claims/name	SAML	user.userprincipalname ...
http://schemas.xmlsoap.org/ws/2005/05/identity/claims/surname	SAML	user.surname ...

∨ Advanced settings

18. For the **Name** field, enter **Email**.

19. Expand **Choose name format**.

- For the **Name format** field, select **Unspecified**.
- Make sure that the **Source** is set to **Attribute**.
- For the **Source attribute** field, choose the drop-down arrow and select **user.mail**.
- Choose **Save**.

Manage claim ... ×

Save
× Discard changes
 🗨️ Got feedback?

Name *

Namespace

∧ Choose name format

Name format

Source * Attribute
 Transformation
 Directory schema extension

Source attribute *

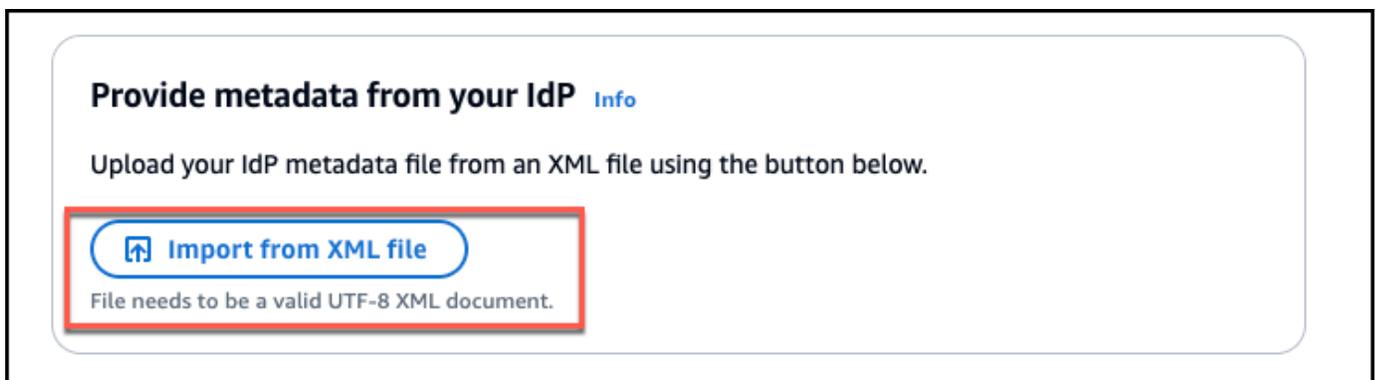
∨ Claim conditions

∨ Advanced SAML claims options

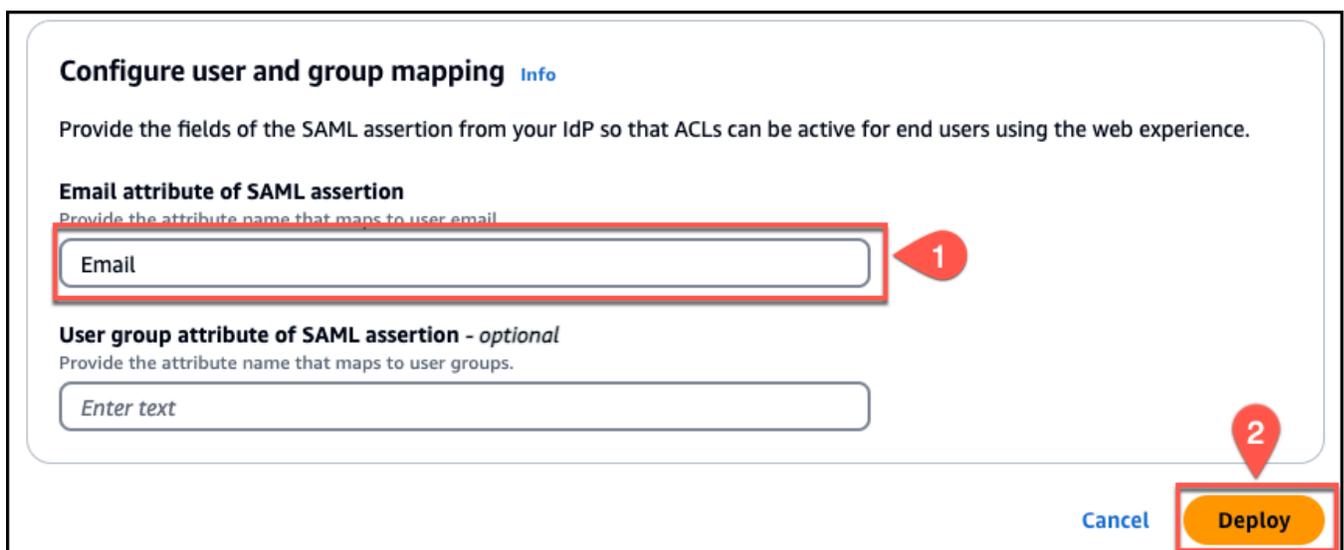
20. Go back to your application page. In the left navigation pane of your application page, choose **Users and groups**.

21. In the **Users** table, select the user that you created earlier. To finish assigning users, choose **Assign**. Continue with the next steps.

- a. If you do not see the user you want to add to your application, choose **+ Add user/group**.
 - b. In the **Add Assignment** page, choose **None Selected**.
 - c. In the right pane, select the user or search for the user in the search bar and then select the user.
 - d. Choose **Select** and then choose **Assign**.
22. In the **Users and groups** page, choose the user name. On the user page, verify that the **User principal name** and **Identities** fields are populated.
 23. Go back to the Amazon Q console, and make sure you're on the **Deploy web experience** page.
 24. Scroll down to the **Provide metadata from your IdP** section. To upload the metadata XML file that you saved in your previous steps, choose **Import from XML**.



25. In the **Configure user and group mapping** section, do the following:
 - For **Email attribute of SAML assertion** – Enter the attribute name that you provided in the Entra ID console. For example, **Email** could be an attribute name.



Note

Make sure there are no spaces at the end of **Email**.

- For [User group field attribute of SAML assertion - optional](#) – Enter an optional user group attribute.

26. Choose **Deploy**.

27. Once deployment finishes, a URL should appear on your Amazon Q application page under **Deployed URL**.

28. Choose the URL to open your Amazon Q web experience and enter credentials for a user that has access to the web experience.

If you encounter HTTP status code 403 (Forbidden) errors, see [Troubleshooting Amazon Q Business and identity provider integration](#).

Setting up Amazon Q Business with Okta as identity provider

The following steps show how to integrate Amazon Q Business with Okta as your SAML 2.0-compliant identity provider (IdP). Integrating Amazon Q with Okta requires that you switch between tasks on the Amazon Q console and the Okta admin console.

Prerequisites

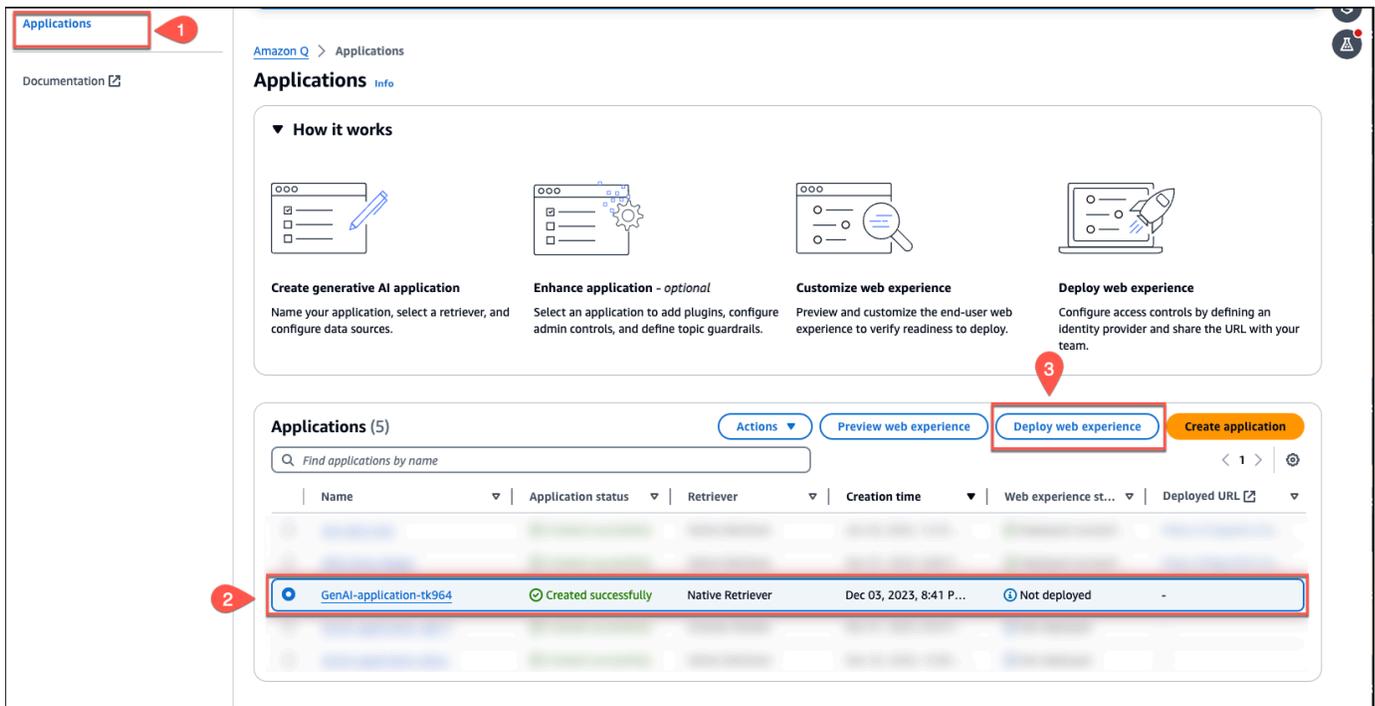
Before you start to integrate Amazon Q with Okta, make sure that you have completed the following tasks:

- Created an Amazon Q Business application, selected a retriever, added your desired data sources, and previewed Amazon Q Business web experience.
- Created an Okta account, added at least one user, assigned users to their groups, and provided each user with a valid email address. For more information, see [Manage users](#) on the *Okta Help Center*.

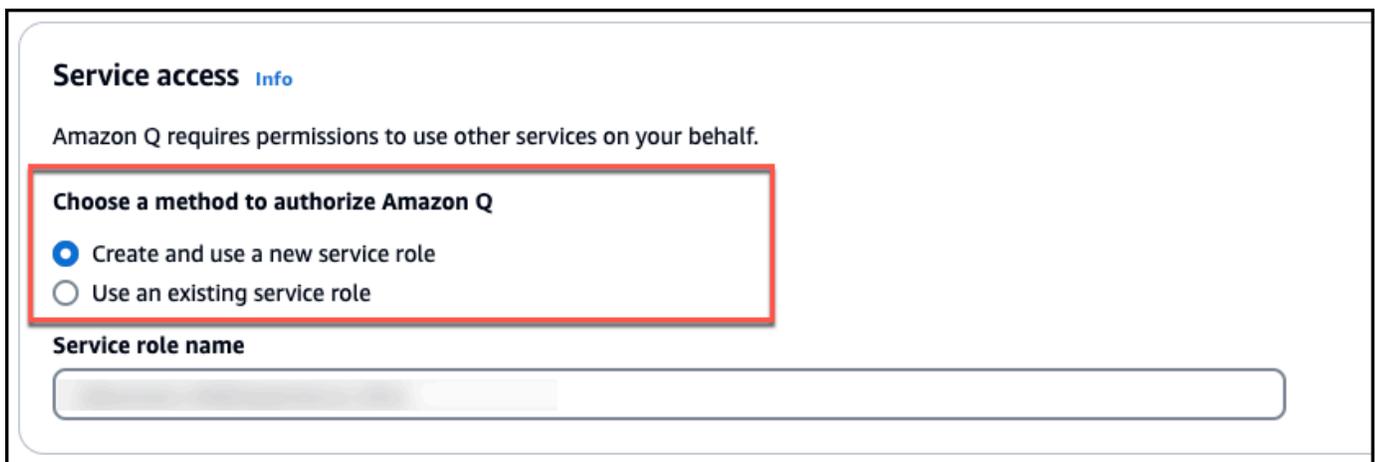
To integrate Amazon Q with Okta

1. In the Amazon Q console, choose your application for integrating with Okta.

- On the **Applications** page, from **Applications**, choose the application you want to deploy. Then, choose **Deploy web experience**.



- On the **Deploy web experience** page, for **Service access**, choose to **Create a use a new service role** or **Use an existing service role**. If you choose to create a new service role, Amazon Q, will automatically create a name for it.



- In the **Configure your [Identity provider](#)** section, do the following:
 - Copy the **[Assertion consumer service\(ACS\) URL](#)** displayed on the console to a text editor of your choice
 - Copy the **[Audience URI \(SP EntityID\)](#)** displayed on the console to a text editor of your choice.

Configure your identity provider (IdP) Info

Configure your SAML support on your IdP. Provide the Assertion Consumer Service (ACS) URL and Audience URI information below to your IdP. This causes IdP metadata to be generated.

Application consumer service (ACS) URL

The ACS URL is the endpoint where the SAML response will be sent.

1

Audience URI (SP Entity ID)

Determines the intended recipient or audience for the SAML assertion.

2

You will use this information later in this procedure.

5. Then, go to the Okta admin console. In the left navigation pane, choose **Applications**, and then choose **Create App Integration**.

Q Search for people, apps and groups

?
☰

- Dashboard ▼
- Directory ▼
- Customizations ▼
- Applications ▲
- Applications
- Self Service
- API Service Integrations
- Security ▼
- Workflow ▼
- Reports ▼
- Settings ▼

Applications

Create App Integration

Browse App Catalog

Assign Users to App

More ▼

STATUS		
ACTIVE	0	<div style="display: flex; align-items: center; margin-bottom: 5px;"> <div style="border: 1px solid #ccc; border-radius: 50%; padding: 2px 5px; margin-right: 5px;"> </div> Okta Admin Console </div>
INACTIVE	0	<div style="display: flex; align-items: center; margin-bottom: 5px;"> <div style="border: 1px solid #ccc; border-radius: 50%; padding: 2px 5px; margin-right: 5px;"> </div> Okta Browser Plugin </div>
		<div style="display: flex; align-items: center; margin-bottom: 5px;"> <div style="border: 1px solid #ccc; border-radius: 50%; padding: 2px 5px; margin-right: 5px;"> </div> Okta Dashboard </div>
		<div style="display: flex; align-items: center; margin-bottom: 5px;"> <div style="border: 1px solid #ccc; border-radius: 50%; padding: 2px 5px; margin-right: 5px;"> </div> Okta Workflows <small>Client ID: 0oa6omk7ofkFuBSP6697</small> </div>
		<div style="display: flex; align-items: center; margin-bottom: 5px;"> <div style="border: 1px solid #ccc; border-radius: 50%; padding: 2px 5px; margin-right: 5px;"> </div> Okta Workflows OAuth <small>Client ID: 0oa6omk7rsVMIOshC697</small> </div>

6. On the **Create a new app integration** page, choose **SAML 2.0** and then choose **Next**.



Create a new app integration

Sign-in method

[Learn More](#)

- OIDC - OpenID Connect**
Token-based OAuth 2.0 authentication for Single Sign-On (SSO) through API endpoints. Recommended if you intend to build a custom app integration with the Okta Sign-In Widget.
- SAML 2.0**
XML-based open standard for SSO. Use if the Identity Provider for your application only supports SAML.
- SWA - Secure Web Authentication**
Okta-specific SSO method. Use if your application doesn't support OIDC or SAML.
- API Services**
Interact with Okta APIs using the scoped OAuth 2.0 access tokens for machine-to-machine authentication.

Cancel

Next

- On the **Create SAML Integration** page, for **General Settings**, in **App name**, enter a name for the application and choose **Next**.

- Dashboard
- Directory
- Customizations
- Applications
- Security
- Workflow
- Reports
- Settings

Create SAML Integration

1 General Settings

2 Configure SAML

3 Feedback

1 General Settings

App name

App name

App logo (optional)



App visibility

Do not display application icon to users

Cancel

Next

- On the **Create SAML Integration** page, for **Configure SAML**, in the **SAML Settings** section, do the following:

- a. For the **Single sign-on URL** field, enter the **Assertion Consumer Service(ACS) URL** that you copied from the Amazon Q console.
- b. For the **Audience URI (SP Entity ID)** field, enter the **Audience URI (SP Entity ID)** that you copied from the Amazon Q console.

Create SAML Integration

1 General Settings

2 Configure SAML

A SAML Settings

General

Single sign-on URL 

Use this for Recipient URL and Destination URL

Audience URI (SP Entity ID) 

Default RelayState 

If no value is set, a blank RelayState is sent

Name ID format 

Application username 

Update application username on

[Show Advanced Settings](#)

9. Scroll down to the **Attribute Statements (optional)** section, and provide the following information. This information will be used by the Amazon Q application to identify the end user's email address.
 - a. For the **Name** field, provide a name for the email attribute, for example Email.
 - b. For the **Name format** field, leave it set to **Unspecified**.
 - c. For the **Value** field, provide a mapping to the attribute by selecting `user.email` from the dropdown list.
 - d. (Optional) To add more attributes, choose **Add another** and provide an attribute name and a value for each user. Make sure to leave the name format set to **Unspecified** for each user.
 - e. Choose **Next**, and then choose **Finish**.
10. From your Okta app page, select the **Assignments** tab.
11. Select **Assign**. To assign users to your Okta app, choose between **Assign to People** and **Assign to Groups**.

[← Back to Applications](#)

App name

Active View Logs Monitor Imports

General Sign On Import **Assignments**

Assign Convert assignments Search... People

Assign to People
Assign to Groups

Groups

Type

01101110
01101111
01101100
01101100
01101101
01101110
01100111

No users found

REPORTS

- Current Assignments
- Recent Unassignments

SELF SERVICE

You need to enable self service for org managed apps before you can use self service for this app.
[Go to self service settings](#)

Requests Disabled

Approval N/A

Edit

12. To finish assigning users, choose **Done**.
13. Go back to the Okta app **Settings** page, and select the **Sign-on** tab.
14. In the **Metadata details** section, to copy the metadata file XML file and save it in .xml format, choose **Copy**.



App name

Active ▾



[View Logs](#) [Monitor Imports](#)

- General
- Sign On**
- Import
- Assignments

Settings

[Edit](#)

Sign on methods

The sign-on method determines how a user signs into and manages their credentials for an application. Some sign-on methods require additional configuration in the 3rd party application.

Application username is determined by the user profile mapping. [Configure profile mapping](#)

SAML 2.0

Default Relay State

Metadata details

Metadata URL <https://trial-8515555.okta.com/app/exk6omnglsw71XUDb697/sso/saml/metadata>

[Copy](#)

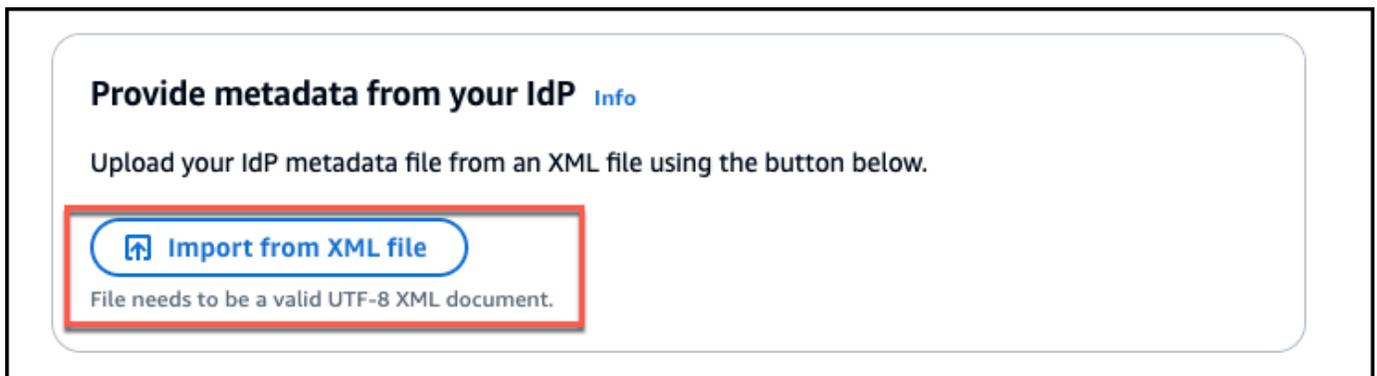
[More details](#)

Note

You can also navigate to the metadata URL and copy the network response payload and paste it in a file that you save in .xml format.

For more information, see [Create SAML app integrations](#) on the *Okta Help Center* website.

15. Go back to the Amazon Q console, and make sure you're on the **Deploy web experience** page.
16. Scroll down to the **Provide metadata from your IdP** section. To upload the metadata XML file that you saved in your previous steps, choose **Import from XML**.



Provide metadata from your IdP [Info](#)

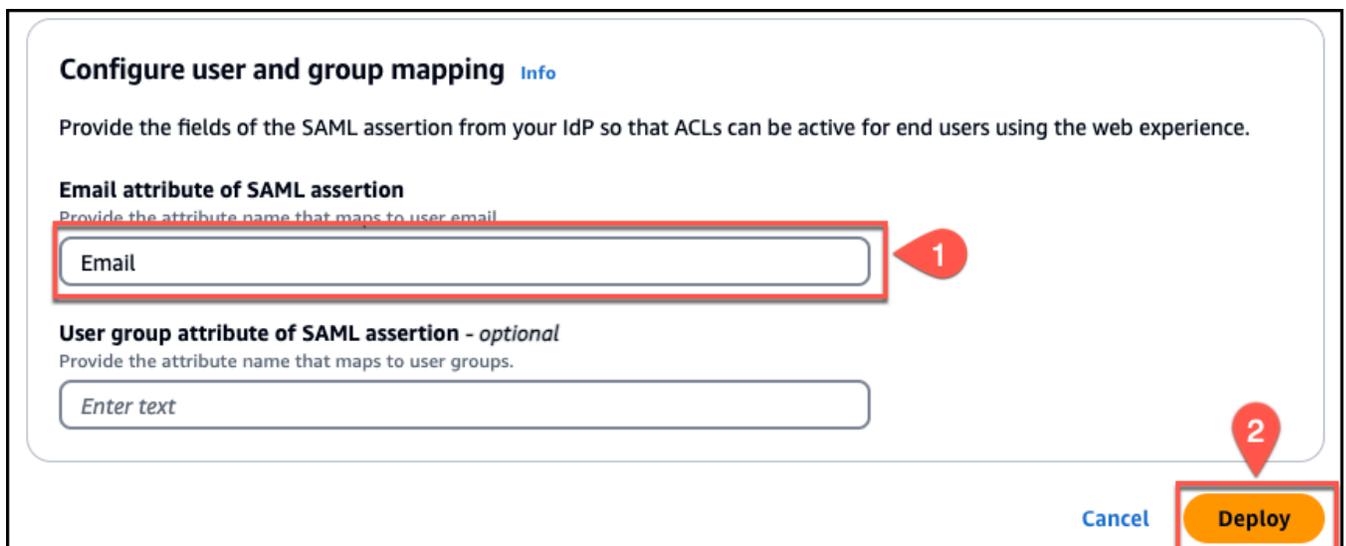
Upload your IdP metadata file from an XML file using the button below.

 **Import from XML file**

File needs to be a valid UTF-8 XML document.

17. In the **Configure user and group mapping** section, do the following:

- For [Email attribute of SAML assertion](#) – Enter the attribute name that you provided in the Entra ID console. For example, **Email** could be an attribute name.



Configure user and group mapping [Info](#)

Provide the fields of the SAML assertion from your IdP so that ACLs can be active for end users using the web experience.

Email attribute of SAML assertion
Provide the attribute name that maps to user email

User group attribute of SAML assertion - optional
Provide the attribute name that maps to user groups.

[Cancel](#) [Deploy](#)

Note

Make sure there are no spaces at the end of **Email**.

- For [User group field attribute of SAML assertion - optional](#) – Enter an optional user group attribute.
18. Choose **Deploy**.
 19. Once deployment finishes, a URL should appear on your Amazon Q application page under **Deployed URL**.
 20. Choose the URL to open your Amazon Q web experience and enter credentials for a user that has access to the web experience.

If you encounter HTTP status code 403 (Forbidden) errors , see [Troubleshooting Amazon Q Business and identity provider integration](#).

Setting up Amazon Q Business with PingIdentity as identity provider

The following steps show how to integrate Amazon Q Business with PingIdentity (Ping) as your SAML 2.0-compliant identity provider (IdP). Integrating Amazon Q with Ping requires that you switch between tasks on the Amazon Q console and your PingIdentity console.

Prerequisites

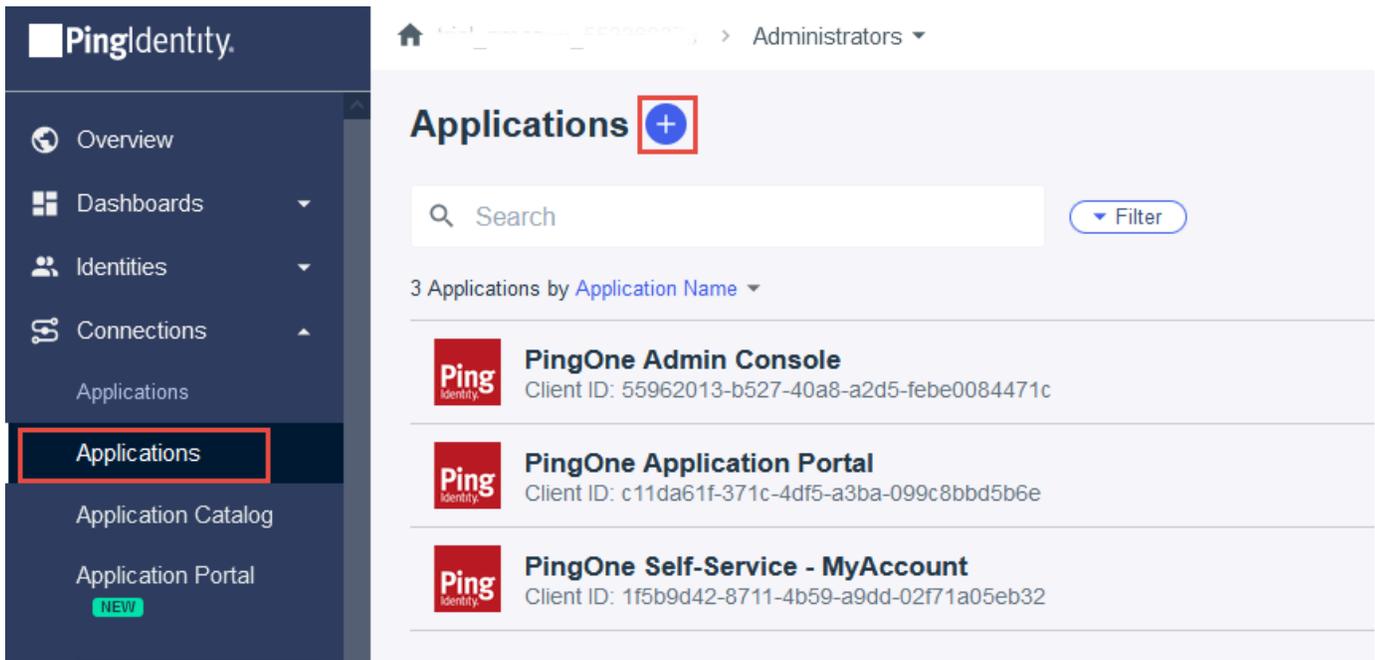
Before you start to integrate Amazon Q with Ping, make sure that you have completed the following tasks:

- Created an Amazon Q Business application, selected a retriever, added your desired data sources, and previewed Amazon Q Business web experience.
- Created a PingIdentity account, added at least one user, and provided each user with a valid email address.

To integrate Amazon Q with Ping

1. In the Amazon Q console, choose your application for integrating with Ping.
2. In the **Application** page, scroll down and choose the **Web experience settings** tab. Choose **Edit**.

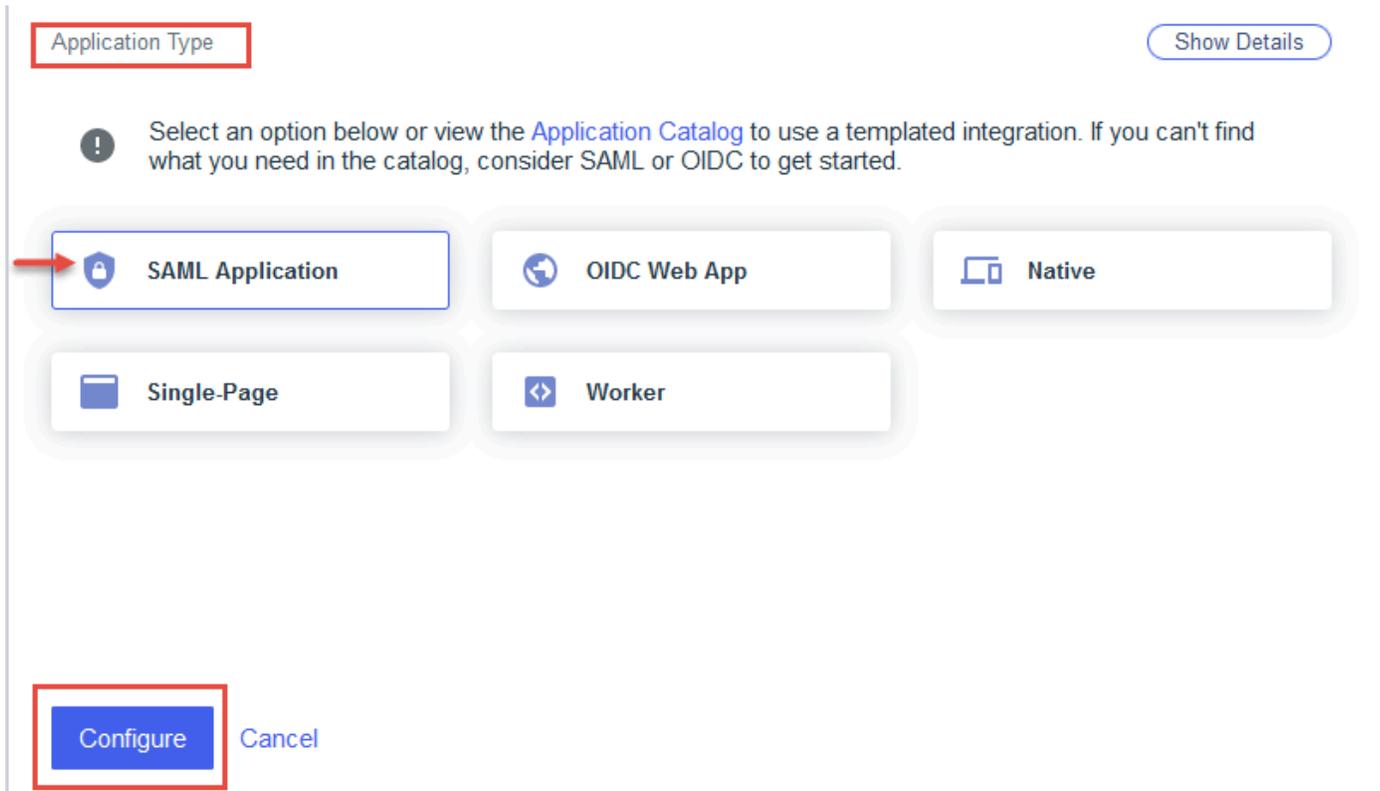
3. For **Service role name**, choose the IAM role that you created for your web experience. Or, choose **Create a new role**. of your Amazon Q application.
4. In the **Configure your identity provider** section, copy the **Assertion Consumer Service (ACS) URL** and the **Audience URI (SP Entity ID)**. You will use them later in this procedure.
5. Go to the PingIdentity console. In the left navigation pane, choose **Applications**.
6. Choose the plus sign (+) next to **Applications** to create a new application.



7. In the **Add Application** section, enter a name for your application and optionally enter a description.



8. In the **Application Type** section, choose **SAML Application** and then choose **Configure**.



9. In the **SAML Configuration** section, choose **Manually Enter** and then do the following:
 - a. For **ACS URLs**, paste the **Application consumer service(ACS) URL** that you copied from the Amazon Q console.
 - b. For **Entity ID**, paste the **Audience URI (SP Identity)** that you copied from the Amazon Q console.
10. Choose **Save**.

Add Application ✕

SAML Configuration

Provide Application Metadata

Import Metadata Import From URL Manually Enter

ACS URLs *

The URL is invalid.

+ Add

Entity ID *

Save Cancel

11. In your application page, choose **Configuration** and then choose **Edit**.

Configuration Overview Attribute Mappings Policies Access

Protocol SAML ⚙️ Attributes 1 Mapped ✎ Policies None Selected ✎ Access All Users ✎ ✎

App Type
Advanced Configuration (SAML)

Description
Not Set

12. Scroll down to the **SUBJECT NAMEID FORMAT** field, set the format to **unspecified**, and then choose **Save**.

The format name will look similar to `urn:oasis:names:tc:SAML:1.1:nameid-format:unspecified`.

Edit Configuration

ENCRYPTION

Enable Encryption

ENTITYID *

SLO ENDPOINT

SUBJECT NAMEID FORMAT
urn:oasis:names:tc:SAML:1.1:nameid-format:unspecified

ASSERTION VALIDITY DURATION (IN SECONDS) *
300

TARGET APPLICATION URL

Enforce Signed AuthnRequest ?

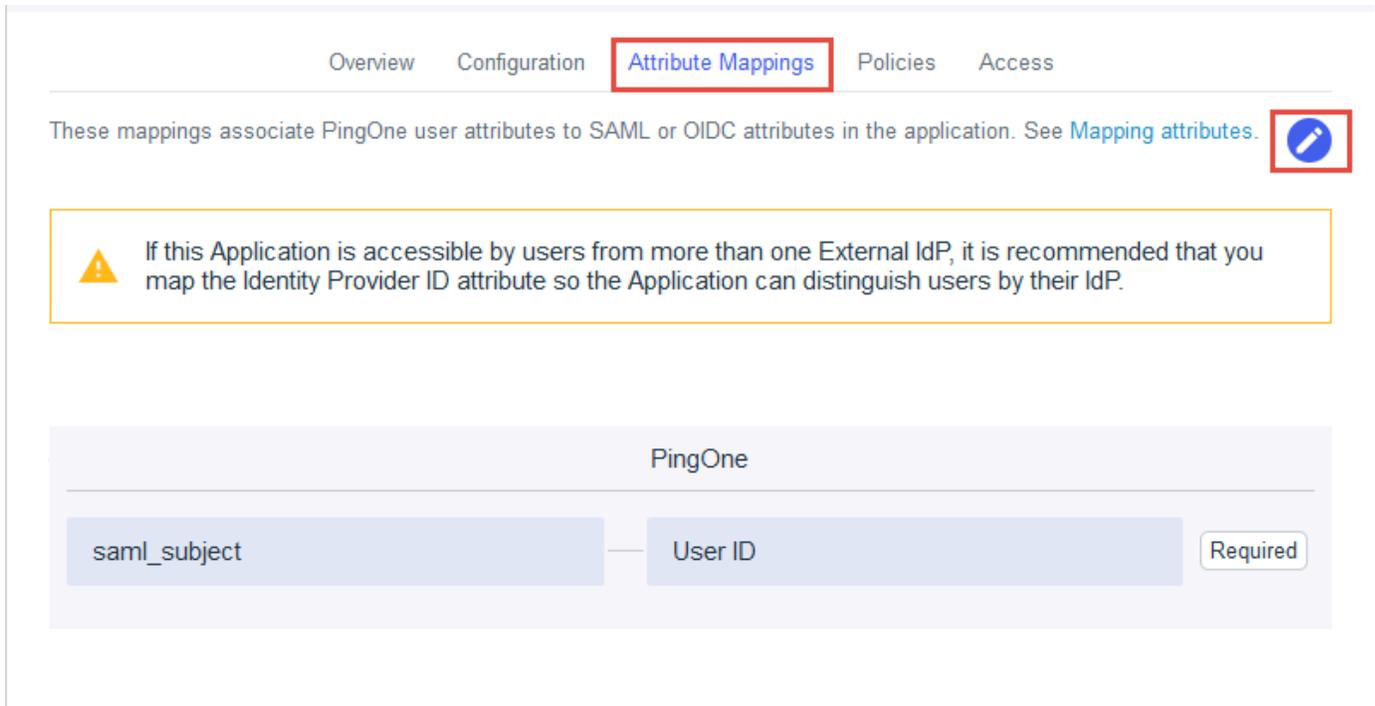
VERIFICATION CERTIFICATE (OPTIONAL)

None Import Choose from list

Select Policy based on RequestedAuthnContext ?

Save Cancel

13. On your application page, choose **Attribute Mappings** and then choose **Edit**.

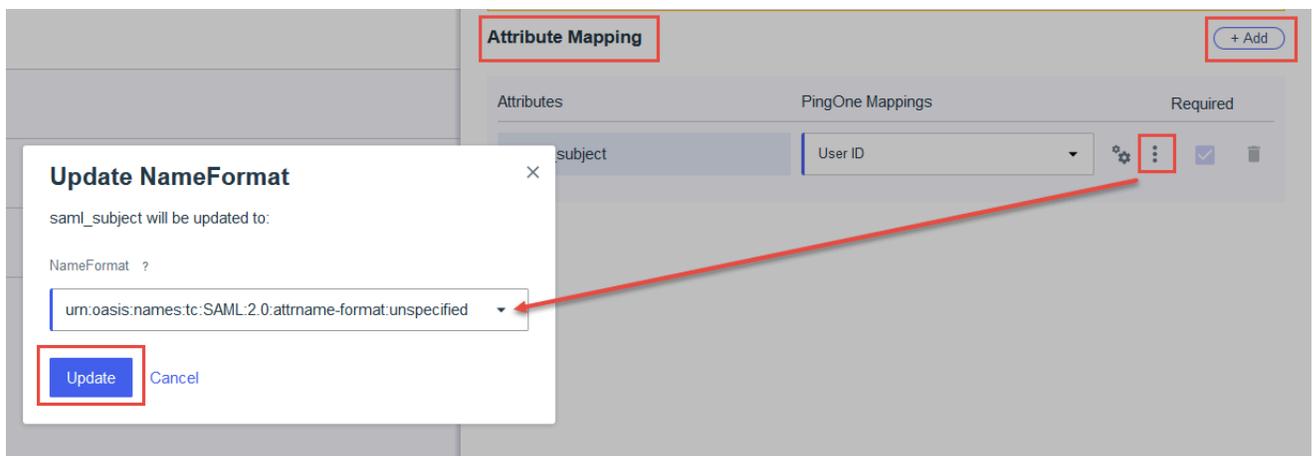


14. On the **Attribute Mapping** page, provide the following information for your application to identify the end user's email address:

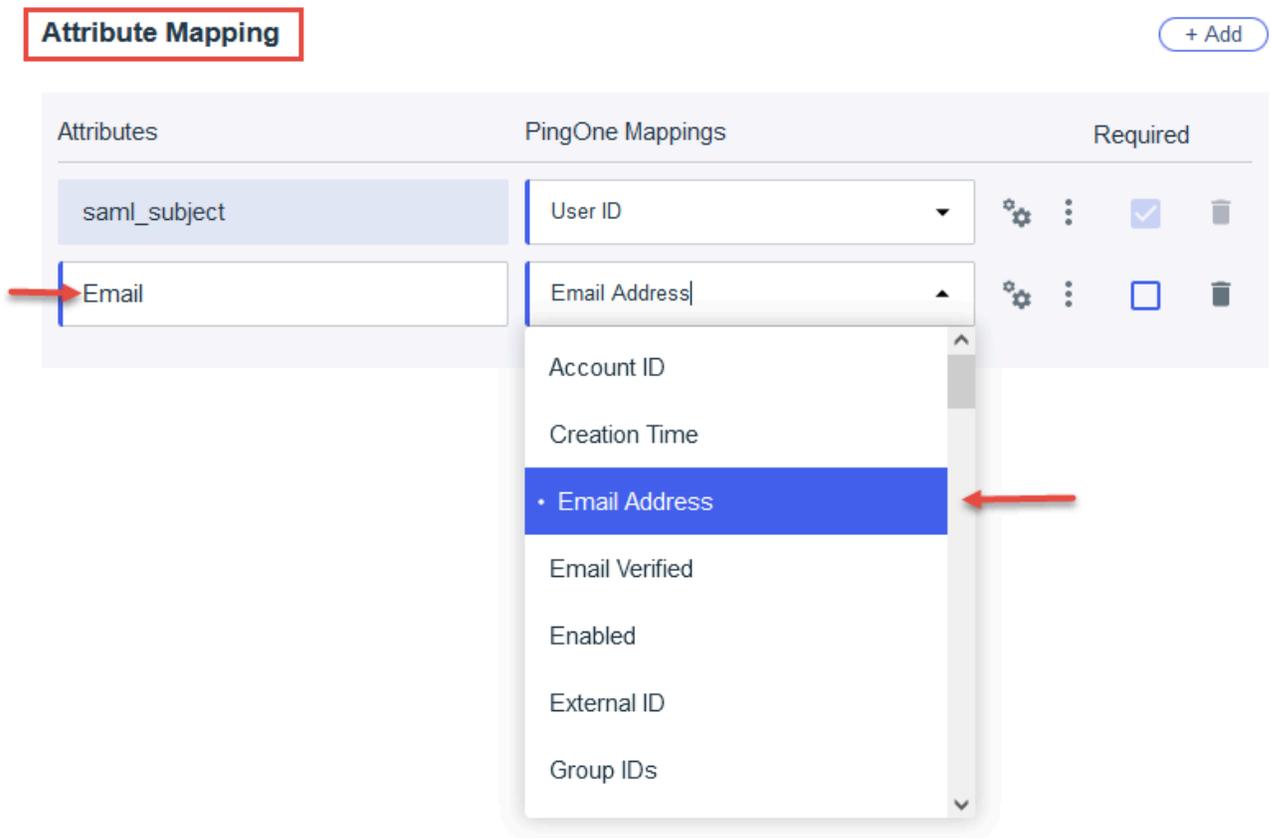
- For the **saml_subject** attribute, leave the **PingOne Mappings** set to **User ID**.
- Choose the update button (three vertical dots), choose **Update NameFormat**, and set the name format to **unspecified**.

The format name will look similar to `urn:oasis:names:tc:SAML:2.0:attrname-format:unspecified`.

- Choose **Update** and then choose **Add**.

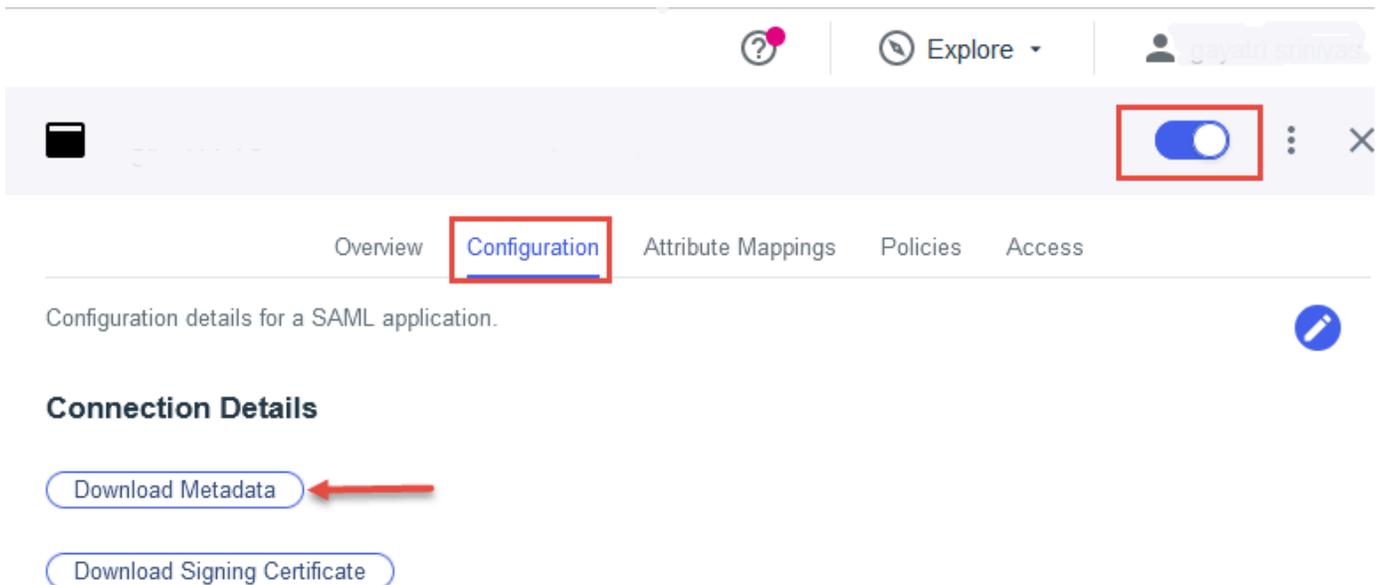


- d. Enter a name for the email attribute, for example, **Email**.
- e. Set the **PingOne Mappings** for email attribute to **Email Address**.
- f. Choose **Save**.

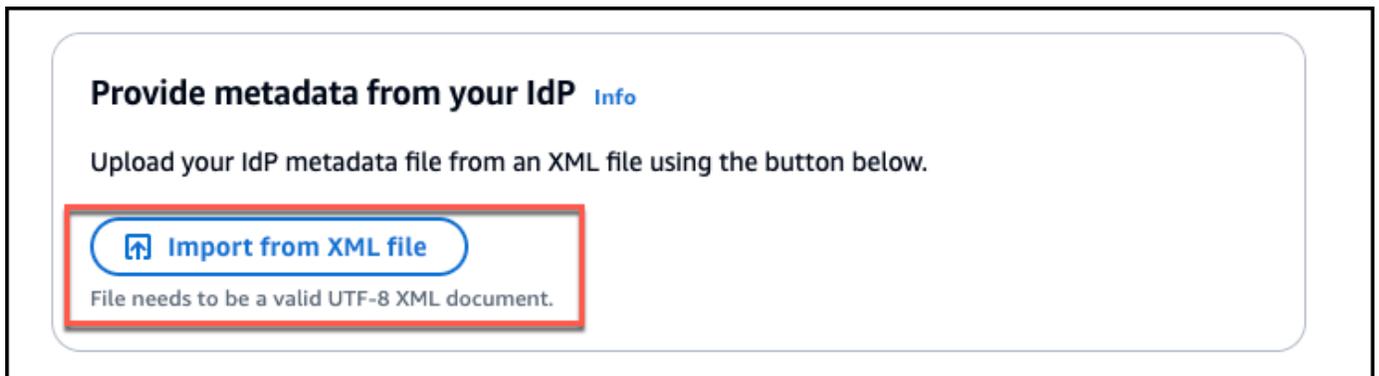


15. Choose **Configuration**. On the **Connection Details** page, choose **Download Metadata**.
16. Choose the enable button next to your application name to enable your application.

By default, all users have access to the application. Choose **Access** if you want to modify the access settings.



17. Go back to the Amazon Q console, and make sure you're on the **Deploy web experience** page.
18. Scroll down to the **Provide metadata from your IdP** section. To upload the metadata XML file that you saved in your previous steps, choose **Import from XML**.



19. In the **Configure user and group mapping** section, do the following:
 - For **Email attribute of SAML assertion** – Enter the attribute name that you provided in the IAM Identity Center console. For example, **Email** could be an attribute name.

Configure user and group mapping [Info](#)

Provide the fields of the SAML assertion from your IdP so that ACLs can be active for end users using the web experience.

Email attribute of SAML assertion
Provide the attribute name that maps to user email

Email

User group attribute of SAML assertion - optional
Provide the attribute name that maps to user groups.

Enter text

Cancel **Deploy**

Note

Make sure there are no spaces at the end of **Emai**l.

- For [User group field attribute of SAML assertion - optional](#) – Enter an optional user group attribute.

20. Choose **Deploy**.

21. Once deployment finishes, a URL should appear on your Amazon Q application page under **Deployed URL**.

22. Choose the URL to open your Amazon Q web experience and enter credentials for a user that has access to the web experience.

If you encounter HTTP status code 403 (Forbidden) errors, see [Troubleshooting Amazon Q Business and identity provider integration](#).

Troubleshooting Amazon Q Business and identity provider integration

This topic helps you troubleshoot issues with opening an Amazon Q Business application after you have integrated Amazon Q with an identity provider.

If you encounter an HTTP status code 403 (Forbidden) error when you open your Amazon Q application, it means that the user is unable to access the application. The following are common causes.

Note

If you're trying to configure end user access to an Amazon Q application through your IdP's application portal instead of a deployed Amazon Q web experience URL, specify the deployed web experience URL as the application start URL in your IdP application settings.

Topics

- [Attribute mappings not set to unspecified](#)
- [Email attribute mismatch](#)
- [User might not have been assigned to the application](#)
- [User's email address is not defined or not mapped correctly](#)
- [Inadequate IAM role permissions](#)

Attribute mappings not set to unspecified

Check the attribute mappings in your identity provider's console. Make sure that the subject attributes and email attributes are set to the **unspecified** format.

For reference, go back to the instructions you followed for integrating Amazon Q with your identity provider:

- For **IAM Identity Center**, see steps 17 and 18 in the [Setting up Amazon Q Business with IAM Identity Center as identity provider](#)
- For **Entra ID**, see steps 18 and 19 in the [Setting up Amazon Q Business with Microsoft Entra ID as identity provider](#)
- For **Okta**, see step 9 in the [Setting up Amazon Q Business with Okta as identity provider](#)
- For **PingIdentity**, see steps 12, 13, and 14 in the [Setting up Amazon Q Business with PingIdentity as identity provider](#)

Email attribute mismatch

You may also get errors because of email attribute name mismatches. Check that the name you entered in the Amazon Q console for **Email attribute** matches the name that you specified in your identity provider attribute mappings page.

For reference, go back to the instructions you followed for integrating Amazon Q with your identity provider:

- For **IAM Identity Center**, see steps 18.b and 22 in the [Setting up Amazon Q Business with IAM Identity Center as identity provider](#)
- For **Entra ID**, see steps 19 and 26 in the [Setting up Amazon Q Business with Microsoft Entra ID as identity provider](#)
- For **Okta**, see steps 9.a and 17 in the [Setting up Amazon Q Business with Okta as identity provider](#)
- For **PingIdentity**, see steps 14.d and 18 in the [Setting up Amazon Q Business with PingIdentity as identity provider](#)

User might not have been assigned to the application

Verify that the user you used to sign in with has access to the web experience. Check the **Assignments** section on your identity provider application page, and confirm that the user is listed and assigned to the web experience.

For reference, go back to the instructions you followed for integrating Amazon Q with your identity provider:

- For **IAM Identity Center**, see step 14 in the [Setting up Amazon Q Business with IAM Identity Center as identity provider](#)
- For **Entra ID**, see steps 21, 22, and 23 in the [Setting up Amazon Q Business with Microsoft Entra ID as identity provider](#)
- For **Okta**, see steps 10 and 11 in the [Setting up Amazon Q Business with Okta as identity provider](#)
- For **PingIdentity**, see step 16 in the [Setting up Amazon Q Business with PingIdentity as identity provider](#)

User's email address is not defined or not mapped correctly

Verify that the user you used to sign in with has a value defined for their email address. Verify that this value is correctly mapped to the email attribute mapping that you configured.

For reference, go back to the instructions you followed for integrating Amazon Q with your identity provider:

- For **IAM Identity Center**, see step 14 in the [Setting up Amazon Q Business with IAM Identity Center as identity provider](#)
- For **Entra ID**, see step 23 in the [Setting up Amazon Q Business with Microsoft Entra ID as identity provider](#)
- For **Okta**, see step 11 in the [Setting up Amazon Q Business with Okta as identity provider](#)
- For **PingIdentity**, see step 16 in the [Setting up Amazon Q Business with PingIdentity as identity provider](#)

Inadequate IAM role permissions

The IAM role used for deploying the Amazon Q web experience might not have the right permissions and trust boundary specified in the policy.

Verify that the IAM role that you've used for granting permissions to the user to access the application has the right service principal listed in the policy.

For reference, see step 8 in the [Steps for deploying your Amazon Q Business web experience](#). If you have created your own IAM role, make sure that the policy provides Amazon Q with permissions to write access relevant Amazon Q API operations. You must also provide a trust policy that allows Amazon Q to assume the role. See [IAM role for an Amazon Q web experience](#) for more information on the policies that you must provide.

Tagging resources

Manage your Amazon Q Business applications and data sources by assigning tags. You can use tags to categorize your Amazon Q resources in various ways. For example, you could categorize by purpose, owner, or application, or any combination. Each tag consists of a *key* and a *value*, both of which you define.

Tags help you to do the following:

- **Identify and organize your AWS resources** – Many AWS services support tagging, so you can assign the same tag to resources in different services to indicate that the resources are related. For example, you can tag an Amazon Kendra retriever and the Amazon Q web experience that uses the retriever with the same tag.
- **Allocate costs** – You activate tags on the AWS Billing and Cost Management dashboard. AWS uses tags to categorize your costs and deliver a monthly cost allocation report to you. For more information, see [Cost Allocation and Tagging](#) in the *AWS Billing User Guide*.
- **Control access to your resources** – You can use tags in AWS Identity and Access Management (IAM) policies that control access to Amazon Q resources. To activate tag-based access control, you can attach these policies to an IAM role or IAM user. For more information, see [Authorization based on tags](#).

You can create and manage tags using the AWS Management Console, the AWS Command Line Interface (AWS CLI), or the Amazon Q API.

Topics

- [Using tags](#)
- [Tag restrictions](#)

Using tags

If you're using the console, you can tag resources when you create them or add them later. You can also use the console to update or remove tags.

If you're using the AWS CLI or the Amazon Q API, use the following operations to manage tags for your resources:

- [CreateApplication](#) – Apply tags when you create an Amazon Q application.
- [CreateDataSource](#) – Apply tags when you create a data source.
- [CreateIndex](#) – Apply tags when you create an Amazon Q retriever and index.
- [CreateRetriever](#) – Apply tags when you create an Amazon Kendra retriever.
- [CreateWebExperience](#) – Apply tags when you create an Amazon Q web experience.
- [CreatePlugin](#) – Apply tags when you create an Amazon Q plugin.
- [ListTagsForResource](#) – View the tags associated with a resource.
- [TagResource](#) – Add and modify tags for a resource.
- [UntagResource](#) – Remove tags from a resource.

Tag restrictions

The following restrictions apply to tags on Amazon Q resources:

- Maximum number of tags – 50
- Maximum key length – 128 characters
- Maximum value length – 256 characters
- Valid characters for key and value – a–z, A–Z, space, and the following characters: _ . : / = + - and @
- Keys and values are case sensitive
- Don't use `aws :` as a prefix for keys; it's reserved for AWS use

Security in Amazon Q Business

Cloud security at AWS is the highest priority. As an AWS customer, you benefit from data centers and network architectures that are built to meet the requirements of the most security-sensitive organizations.

Security is a shared responsibility between AWS and you. The [shared responsibility model](#) describes this as security *of* the cloud and security *in* the cloud:

- **Security of the cloud** – AWS is responsible for protecting the infrastructure that runs AWS services in the AWS Cloud. AWS also provides you with services that you can use securely. Third-party auditors regularly test and verify the effectiveness of our security as part of the [AWS Compliance Programs](#). To learn about the compliance programs that apply to Amazon Q Business, see [AWS Services in Scope by Compliance Program](#).
- **Security in the cloud** – Your responsibility is determined by the AWS service that you use. You are also responsible for other factors including the sensitivity of your data, your company's requirements, and applicable laws and regulations.

This documentation helps you understand how to apply the shared responsibility model when using Amazon Q. The following topics show you how to configure Amazon Q to meet your security and compliance objectives. You also learn how to use other AWS services that help you to monitor and secure your Amazon Q resources.

Topics

- [Data protection in Amazon Q Business](#)
- [Identity and access management for Amazon Q Business](#)
- [Compliance validation for Amazon Q Business](#)
- [Resilience in Amazon Q Business](#)
- [Infrastructure security in Amazon Q Business](#)
- [Cross-service confused deputy prevention](#)
- [Configuration and vulnerability analysis in AWS Identity and Access Management](#)
- [Security best practices](#)

Data protection in Amazon Q Business

The AWS [shared responsibility model](#) applies to data protection in Amazon Q Business. As described in this model, AWS is responsible for protecting the global infrastructure that runs all of the AWS Cloud. You are responsible for maintaining control over your content that is hosted on this infrastructure. You are also responsible for the security configuration and management tasks for the AWS services that you use. For more information about data privacy, see the [Data Privacy FAQ](#). For information about data protection in Europe, see the [AWS Shared Responsibility Model and GDPR](#) blog post on the *AWS Security Blog*.

For data protection purposes, we recommend that you protect AWS account credentials and set up individual users with AWS IAM Identity Center or AWS Identity and Access Management (IAM). That way, each user is given only the permissions necessary to fulfill their job duties. We also recommend that you secure your data in the following ways:

- Use multi-factor authentication (MFA) with each account.
- Use SSL/TLS to communicate with AWS resources. We require TLS 1.2 and recommend TLS 1.3.
- Set up API and user activity logging with AWS CloudTrail.
- Use AWS encryption solutions, along with all default security controls within AWS services.
- Use advanced managed security services such as Amazon Macie, which assists in discovering and securing sensitive data that is stored in Amazon S3.
- If you require FIPS 140-2 validated cryptographic modules when accessing AWS through a command line interface or an API, use a FIPS endpoint. For more information about the available FIPS endpoints, see [Federal Information Processing Standard \(FIPS\) 140-2](#).

We strongly recommend that you never put confidential or sensitive information, such as your customers' email addresses, into tags or free-form text fields such as a **Name** field. This includes when you work with Amazon Q or other AWS services using the console, API, AWS CLI, or AWS SDKs. Any data that you enter into tags or free-form text fields used for names may be used for billing or diagnostic logs. If you provide a URL to an external server, we strongly recommend that you do not include credentials information in the URL to validate your request to that server.

Topics

- [Data encryption](#)
- [Key management](#)

Data encryption

Amazon Q Business supports encryption at rest using an AWS KMS key that's owned by AWS. Amazon Q also uses HTTPS protocol for data in transit.

Important

Amazon Q does not support asymmetric KMS keys. For more information, see [Using Symmetric and Asymmetric Keys](#) in the *AWS Key Management Service Developer Guide*.

Topics

- [Encryption at rest](#)
- [Encryption in transit](#)

Encryption at rest

Amazon Q Business provides encryption by default to protect sensitive customer data at rest using AWS owned encryption keys. Sensitive customer data includes both questions and answers in the Amazon Q web experience and the documents uploaded to Amazon Q index.

The Amazon Q uses the questions and answers to know the conversation context and to provide you with the best answer. The conversation data is automatically removed once the conversation is deleted or is inactive. For more information, see [Conversation management](#). The uploaded documents are used by Amazon Q to retrieve them at runtime to answer your questions.

- **AWS owned keys** – Amazon Q uses these keys by default to automatically encrypt sensitive customer data. You can't view, manage, or use AWS owned keys, or audit their use. However, you don't have to take any action or change any programs to protect the keys that encrypt your data. For more information, see [AWS owned keys](#) in the *AWS Key Management Service Developer Guide*.

Encryption of data at rest by default helps reduce the operational overhead and complexity involved in protecting sensitive data. At the same time, it enables you to build secure applications that meet strict encryption compliance and regulatory requirements.

While you can't disable this layer of encryption or select an alternate encryption type, you can add a second layer of encryption over the existing AWS owned encryption keys by choosing a customer managed key when you create your resources:

- **Customer managed keys (CMK)** – Amazon Q supports the use of a symmetric customer managed key that you create, own, and manage to add a second layer of encryption over the existing AWS owned encryption. If you choose to use a customer managed key, you must provision at least 10 index storage units when you [create an Amazon Q retriever](#).

⚠ Important

Amazon Q does not support asymmetric KMS keys. For more information, see [Using Symmetric and Asymmetric Keys](#) in the *AWS Key Management Service Developer Guide*.

Because you have full control of this layer of encryption, you can perform such tasks as:

- Establishing and maintaining key policies
- Establishing and maintaining IAM policies and grants
- Enabling and disabling key policies
- Rotating key cryptographic material
- Adding tags
- Creating key aliases
- Scheduling keys for deletion

For more information, see [customer managed key](#) in the *AWS Key Management Service Developer Guide*.

📘 Note

If you have created your Amazon Q application using AWS KMS and then you want to migrate to using customer managed key (CMK), you will have to re-create your application.

Topics

- [How Amazon Q Business uses grants in AWS KMS](#)
- [Create a customer managed key \(CMK\)](#)
- [Specifying customer managed key for Amazon Q](#)
- [Monitoring your encryption keys for Amazon Q](#)

How Amazon Q Business uses grants in AWS KMS

Amazon Q requires a [grant](#) to use your customer managed key. When you create a Amazon Q application resource encrypted with a customer managed key, Amazon Q creates a grant on your behalf by sending a [CreateGrant](#) request to AWS KMS. Grants in AWS KMS are used to give Amazon Q access to a KMS key in a customer account.

Amazon Q requires the grant to use your customer managed key for the following internal operations:

- Send [DescribeKey](#) requests to AWS KMS to verify that the symmetric customer managed key ID entered when creating application is valid.
- Send [GenerateDataKeyWithoutPlainText](#) requests to AWS KMS to generate data keys encrypted by your customer managed key.
- Send [Decrypt](#) requests to AWS KMS to decrypt the encrypted data keys so that they can be used to encrypt your data.

You can revoke access to the grant, or remove the service's access to the customer managed key at any time. If you do, Amazon Q won't be able to access any of the data encrypted by the customer managed key, which affects operations that are dependent on that data.

Create a customer managed key (CMK)

You can create a symmetric customer managed key by using the AWS Management Console, or the AWS KMS APIs.

Important

Amazon Q does not support asymmetric KMS keys. For more information, see [Using Symmetric and Asymmetric Keys](#) in the *AWS Key Management Service Developer Guide*.

To create a symmetric customer managed key

Follow the steps for [Creating symmetric customer managed key](#) in the *AWS Key Management Service Developer Guide*.

Key policy

Key policies control access to your customer managed key. Every customer managed key must have exactly one key policy, which contains statements that determine who can use the key and how they can use it. When you create your customer managed key, you can specify a key policy. For more information, see [Managing access to customer managed keys](#) in the *AWS Key Management Service Developer Guide*.

To use your customer managed key with your Amazon Q resources, the following API operations must be permitted in the key policy:

- [kms:CreateGrant](#) – Adds a grant to a customer managed key. Grants control access to a specified KMS key, which allows access to [grant operation](#) Amazon Q requires. For more information about [Using Grants](#), see the *AWS Key Management Service Developer Guide*.

This allows Amazon Q to do the following:

- Call `GenerateDataKeyWithoutPlainText` to generate an encrypted data key and store it, because the data key isn't immediately used to encrypt.
- Call `Decrypt` to use the stored encrypted data key to access encrypted data.
- Set up a retiring principal to allow the service to `RetireGrant`.
- [kms:DescribeKey](#) – Provides the customer managed key details to allow Amazon Q to validate the key.

The following are policy statement examples you can add for Amazon Q

```
"Statement": [{
  "Sid": "Allow access to principals authorized to use Amazon Q",
  "Effect": "Allow",
  "Principal": {
    "AWS": "*"
  },
  "Action": [
    "kms:DescribeKey",
    "kms:CreateGrant"
  ],
  "Resource": "*",
  "Condition": {
    "StringEquals": {
      "kms:ViaService": "qbusiness.region.amazonaws.com",
      "kms:CallerAccount": "111122223333"
    }
  }
}]
```

```

    }
  },
  {
    "Sid": "Allow access for key administrators",
    "Effect": "Allow",
    "Principal": {
      "AWS": "arn:aws:iam::111122223333:root"
    },
    "Action": [
      "kms:*"
    ],
    "Resource": "arn:aws:kms:region:111122223333:key/key_ID"
  },
  {
    "Sid": "Allow read-only access to key metadata to the account",
    "Effect": "Allow",
    "Principal": {
      "AWS": "arn:aws:iam::111122223333:root"
    },
    "Action": [
      "kms:Describe*",
      "kms:Get*",
      "kms:List*",
      "kms:RevokeGrant"
    ],
    "Resource": "*"
  }
]

```

For more information about [specifying permissions in a policy](#), see the *AWS Key Management Service Developer Guide*.

For more information about [troubleshooting key access](#), see the *AWS Key Management Service Developer Guide*.

Specifying customer managed key for Amazon Q

You can specify a customer managed key as a second layer encryption for your Amazon Q application resource.

When you create your application, you can specify the data key by entering a **KMS ID**, which Amazon Q uses to encrypt the identifiable personal data stored by the application.

KMS ID – A [key identifier](#) for an AWS KMS customer managed key. Enter a key ID, key ARN, alias name, or alias ARN.

Any resources you create under your Amazon Q application will be encrypted with the same key.

Monitoring your encryption keys for Amazon Q

When you use an AWS KMS customer managed key with your Amazon Q resources, you can use [AWS CloudTrail](#) or [Amazon CloudWatch Logs](#) to track requests that Amazon Q sends to AWS KMS.

The following examples are AWS CloudTrail events for CreateGrant, GenerateDataKey, Decrypt, and DescribeKey to monitor KMS operations called by Amazon Q to access data encrypted by your customer managed key.

CreateGrant

When you use an AWS KMS customer managed key to encrypt your application, Amazon Q sends a CreateGrant request on your behalf to access the KMS key in your AWS account. The grant that Amazon Q creates are specific to the resource associated with the AWS KMS customer managed key. In addition, Amazon Q uses the RetireGrant operation to remove a grant when you delete a resource.

The following example event records the CreateGrant operation:

```
{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "AssumedRole",
    "principalId": "AROAIQDTESTANDEXAMPLE:Sampleuser01",
    "arn": "arn:aws:sts::111122223333:assumed-role/Admin/Sampleuser01",
    "accountId": "111122223333",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE3",
    "sessionContext": {
      "sessionIssuer": {
        "type": "Role",
        "principalId": "AROAIQDTESTANDEXAMPLE:Sampleuser01",
        "arn": "arn:aws:sts::111122223333:assumed-role/Admin/Sampleuser01",
        "accountId": "111122223333",
        "userName": "Admin"
      },
      "webIdFederationData": {},
      "attributes": {
```



```

        "mfaAuthenticated": "false",
        "creationDate": "2021-04-22T17:02:00Z"
    }
},
"invokedBy": "qbusiness.amazonaws.com"
},
"eventTime": "2021-04-22T17:07:02Z",
"eventSource": "kms.amazonaws.com",
"eventName": "CreateGrant",
"awsRegion": "us-west-2",
"sourceIPAddress": "172.12.34.56",
"userAgent": "ExampleDesktop/1.0 (V1; OS)",
"requestParameters": {
    "retiringPrincipal": "qbusiness.region.amazonaws.com",
    "operations": [
        "CreateGrant",
        "RetireGrant",
        "GenerateDataKey",
        "GenerateDataKeyWithoutPlaintext",
        "Encrypt",
        "ReEncryptTo",
        "ReEncryptFrom",
        "Decrypt",
        "DescribeKey"
    ],
    "keyId": "arn:aws:kms:us-
west-2:111122223333:key/1234abcd-12ab-34cd-56ef-123456SAMPLE",
    "granteePrincipal": "qbusiness.region.amazonaws.com"
},
"responseElements": {
    "grantId":
"0ab0ac0d0b000f00ea00cc0a0e00fc00bce000c000f0000000c0bc0a0000aaafSAMPLE"
},
"requestID": "ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
"eventID": "ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
"readOnly": false,
"resources": [
    {
        "accountId": "111122223333",
        "type": "AWS::KMS::Key",
        "ARN": "arn:aws:kms:us-
west-2:111122223333:key/1234abcd-12ab-34cd-56ef-123456SAMPLE"
    }
],

```

```

    "eventType": "AwsApiCall",
    "managementEvent": true,
    "eventCategory": "Management",
    "recipientAccountId": "111122223333"
  }

```

GenerateDataKey

When you use an AWS KMS customer managed key for your application, Amazon Q creates a unique table key. It sends a `GenerateDataKey` request to AWS KMS that specifies the AWS KMS customer managed key for the application.

The following example event records the `GenerateDataKey` operation:

```

{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "AWSService",
    "invokedBy": "qbusiness.amazonaws.com"
  },
  "eventTime": "2023-11-24T01:50:25Z",
  "eventSource": "kms.amazonaws.com",
  "eventName": "GenerateDataKey",
  "awsRegion": "us-west-2",
  "sourceIPAddress": "172.12.34.56",
  "userAgent": "ExampleDesktop/1.0 (V1; OS)",
  "requestParameters": {
    "keyId": "arn:aws:kms:us-west-2:398547360552:key/ba6c9092-ad4d-41c3-937a-f02177ae147e",
    "keySpec": "AES_256"
  },
  "responseElements": null,
  "requestID": "4bd8e018-90d0-4b93-bc8d-32338578a158",
  "eventID": "aca6cb5b-44bb-3ed6-afdd-736432323356",
  "readOnly": true,
  "resources": [
    {
      "accountId": "111122223333",
      "type": "AWS::KMS::Key",
      "ARN": "arn:aws:kms:us-west-2:398547360552:key/ba6c9092-ad4d-41c3-937a-f02177ae147e"
    }
  ],
}

```

```

    "eventType": "AwsApiCall",
    "managementEvent": true,
    "recipientAccountId": "398547360552",
    "sharedEventID": "57393866-c398-4fd6-a259-d6cb001c7cf9",
    "eventCategory": "Management"
  }

```

Decrypt

When you access an encrypted application, Amazon Q calls the Decrypt operation to use the stored encrypted data key to access the encrypted data.

The following example event records the Decrypt operation.

```

{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "AWSService",
    "invokedBy": "qbusiness.amazonaws.com"
  },
  "eventTime": "2021-04-22T17:10:51Z",
  "eventSource": "kms.amazonaws.com",
  "eventName": "Decrypt",
  "awsRegion": "us-west-2",
  "sourceIPAddress": "172.12.34.56",
  "userAgent": "ExampleDesktop/1.0 (V1; OS)",
  "requestParameters": {
    "keyId": "arn:aws:kms:us-west-2:111122223333:key/1234abcd-12ab-34cd-56ef-123456SAMPLE",
    "encryptionAlgorithm": "SYMMETRIC_DEFAULT"
  },
  "responseElements": null,
  "requestID": "ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
  "eventID": "ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
  "readOnly": true,
  "resources": [
    {
      "accountId": "111122223333",
      "type": "AWS::KMS::Key",
      "ARN": "arn:aws:kms:us-west-2:111122223333:key/1234abcd-12ab-34cd-56ef-123456SAMPLE"
    }
  ],
  "eventType": "AwsApiCall",

```

```

    "managementEvent": true,
    "eventCategory": "Management",
    "recipientAccountId": "111122223333",
    "sharedEventID": "dc129381-1d94-49bd-b522-f56a3482d088"
  }

```

DescribeKey

Amazon Q uses the DescribeKey operation to verify if the AWS KMS customer managed key associated with your application exists in the account and region.

The following example event records DescribeKey operation:

```

{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "AssumedRole",
    "principalId": "AROAIQDTESTANDEXAMPLE:Sampleuser01",
    "arn": "arn:aws:sts::111122223333:assumed-role/Admin/Sampleuser01",
    "accountId": "111122223333",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE3",
    "sessionContext": {
      "sessionIssuer": {
        "type": "Role",
        "principalId": "AROAIQDTESTANDEXAMPLE:Sampleuser01",
        "arn": "arn:aws:sts::111122223333:assumed-role/Admin/
Sampleuser01",
        "accountId": "111122223333",
        "userName": "Admin"
      },
      "webIdFederationData": {},
      "attributes": {
        "mfaAuthenticated": "false",
        "creationDate": "2021-04-22T17:02:00Z"
      }
    },
    "invokedBy": "qbusiness.amazonaws.com"
  },
  "eventTime": "2021-04-22T17:07:02Z",
  "eventSource": "kms.amazonaws.com",
  "eventName": "DescribeKey",
  "awsRegion": "us-west-2",
  "sourceIPAddress": "172.12.34.56",
  "userAgent": "ExampleDesktop/1.0 (V1; OS)",

```

```
"requestParameters": {
  "keyId": "00dd0db0-0000-0000-ac00-b0c000SAMPLE"
},
"responseElements": null,
"requestID": "ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
"eventID": "ff000af-00eb-00ce-0e00-ea000fb0fba0SAMPLE",
"readOnly": true,
"resources": [
  {
    "accountId": "111122223333",
    "type": "AWS::KMS::Key",
    "ARN": "arn:aws:kms:us-
west-2:111122223333:key/1234abcd-12ab-34cd-56ef-123456SAMPLE"
  }
],
"eventType": "AwsApiCall",
"managementEvent": true,
"eventCategory": "Management",
"recipientAccountId": "111122223333"
}
```

Encryption in transit

Amazon Q Business uses the HTTPS protocol to communicate with your client application. It uses HTTPS and AWS signatures to communicate with other services on your application's behalf. .

Key management

Amazon Q Business encrypts the contents of your index using one of two types of keys. You can choose one of the following:

- An AWS-owned AWS KMS. This is the default.
- A customer-managed KMS key. You can create the key when you are creating an Amazon Q application, retriever, index, web experience, data source, or plugins, or you can create the key using the AWS KMS console. Select a symmetric encryption customer-managed KMS key.

Important

Amazon Q does not support asymmetric KMS keys. For more information, see [Using Symmetric and Asymmetric Keys](#) in the *AWS Key Management Service Developer Guide*.

Identity and access management for Amazon Q Business

AWS Identity and Access Management (IAM) is an AWS service that helps an administrator securely control access to AWS resources. IAM administrators control who can be *authenticated* (signed in) and *authorized* (have permissions) to use Amazon Q resources. IAM is an AWS service that you can use with no additional charge.

Topics

- [Audience](#)
- [Authenticating with identities](#)
- [Managing access using policies](#)
- [How Amazon Q Business works with IAM](#)
- [Identity-based policy examples for Amazon Q Business](#)
- [Troubleshooting Amazon Q Business identity and access](#)

Audience

How you use AWS Identity and Access Management (IAM) differs, depending on the work that you do in Amazon Q.

Service user – If you use the Amazon Q service to do your job, then your administrator provides you with the credentials and permissions that you need. As you use more Amazon Q features to do your work, you might need additional permissions. Understanding how access is managed can help you request the right permissions from your administrator. If you cannot access a feature in Amazon Q, see [Troubleshooting Amazon Q Business identity and access](#).

Service administrator – If you're in charge of Amazon Q resources at your company, you probably have full access to Amazon Q. It's your job to determine which Amazon Q features and resources your service users should access. You must then submit requests to your IAM administrator to change the permissions of your service users. Review the information on this page to understand the basic concepts of IAM. To learn more about how your company can use IAM with Amazon Q, see [How Amazon Q Business works with IAM](#).

IAM administrator – If you're an IAM administrator, you might want to learn details about how you can write policies to manage access to Amazon Q. To view example Amazon Q identity-based policies that you can use in IAM, see [Identity-based policy examples for Amazon Q Business](#).

Authenticating with identities

Authentication is how you sign in to AWS using your identity credentials. You must be *authenticated* (signed in to AWS) as the AWS account root user, as an IAM user, or by assuming an IAM role.

You can sign in to AWS as a federated identity by using credentials provided through an identity source. AWS IAM Identity Center (IAM Identity Center) users, your company's single sign-on authentication, and your Google or Facebook credentials are examples of federated identities. When you sign in as a federated identity, your administrator previously set up identity federation using IAM roles. When you access AWS by using federation, you are indirectly assuming a role.

Depending on the type of user you are, you can sign in to the AWS Management Console or the AWS access portal. For more information about signing in to AWS, see [How to sign in to your AWS account](#) in the *AWS Sign-In User Guide*.

If you access AWS programmatically, AWS provides a software development kit (SDK) and a command line interface (CLI) to cryptographically sign your requests by using your credentials. If you don't use AWS tools, you must sign requests yourself. For more information about using the recommended method to sign requests yourself, see [Signing AWS API requests](#) in the *IAM User Guide*.

Regardless of the authentication method that you use, you might be required to provide additional security information. For example, AWS recommends that you use multi-factor authentication (MFA) to increase the security of your account. To learn more, see [Multi-factor authentication](#) in the *AWS IAM Identity Center User Guide* and [Using multi-factor authentication \(MFA\) in AWS](#) in the *IAM User Guide*.

AWS account root user

When you create an AWS account, you begin with one sign-in identity that has complete access to all AWS services and resources in the account. This identity is called the AWS account *root user* and is accessed by signing in with the email address and password that you used to create the account. We strongly recommend that you don't use the root user for your everyday tasks. Safeguard your root user credentials and use them to perform the tasks that only the root user can perform. For the complete list of tasks that require you to sign in as the root user, see [Tasks that require root user credentials](#) in the *IAM User Guide*.

Federated identity

As a best practice, require human users, including users that require administrator access, to use federation with an identity provider to access AWS services by using temporary credentials.

A *federated identity* is a user from your enterprise user directory, a web identity provider, the AWS Directory Service, the Identity Center directory, or any user that accesses AWS services by using credentials provided through an identity source. When federated identities access AWS accounts, they assume roles, and the roles provide temporary credentials.

For centralized access management, we recommend that you use AWS IAM Identity Center. You can create users and groups in IAM Identity Center, or you can connect and synchronize to a set of users and groups in your own identity source for use across all your AWS accounts and applications. For information about IAM Identity Center, see [What is IAM Identity Center?](#) in the *AWS IAM Identity Center User Guide*.

IAM users and groups

An [IAM user](#) is an identity within your AWS account that has specific permissions for a single person or application. Where possible, we recommend relying on temporary credentials instead of creating IAM users who have long-term credentials such as passwords and access keys. However, if you have specific use cases that require long-term credentials with IAM users, we recommend that you rotate access keys. For more information, see [Rotate access keys regularly for use cases that require long-term credentials](#) in the *IAM User Guide*.

An [IAM group](#) is an identity that specifies a collection of IAM users. You can't sign in as a group. You can use groups to specify permissions for multiple users at a time. Groups make permissions easier to manage for large sets of users. For example, you could have a group named *IAMAdmins* and give that group permissions to administer IAM resources.

Users are different from roles. A user is uniquely associated with one person or application, but a role is intended to be assumable by anyone who needs it. Users have permanent long-term credentials, but roles provide temporary credentials. To learn more, see [When to create an IAM user \(instead of a role\)](#) in the *IAM User Guide*.

IAM roles

An [IAM role](#) is an identity within your AWS account that has specific permissions. It is similar to an IAM user, but is not associated with a specific person. You can temporarily assume an IAM role in

the AWS Management Console by [switching roles](#). You can assume a role by calling an AWS CLI or AWS API operation or by using a custom URL. For more information about methods for using roles, see [Using IAM roles](#) in the *IAM User Guide*.

IAM roles with temporary credentials are useful in the following situations:

- **Federated user access** – To assign permissions to a federated identity, you create a role and define permissions for the role. When a federated identity authenticates, the identity is associated with the role and is granted the permissions that are defined by the role. For information about roles for federation, see [Creating a role for a third-party Identity Provider](#) in the *IAM User Guide*. If you use IAM Identity Center, you configure a permission set. To control what your identities can access after they authenticate, IAM Identity Center correlates the permission set to a role in IAM. For information about permissions sets, see [Permission sets](#) in the *AWS IAM Identity Center User Guide*.
- **Temporary IAM user permissions** – An IAM user or role can assume an IAM role to temporarily take on different permissions for a specific task.
- **Cross-account access** – You can use an IAM role to allow someone (a trusted principal) in a different account to access resources in your account. Roles are the primary way to grant cross-account access. However, with some AWS services, you can attach a policy directly to a resource (instead of using a role as a proxy). To learn the difference between roles and resource-based policies for cross-account access, see [How IAM roles differ from resource-based policies](#) in the *IAM User Guide*.
- **Cross-service access** – Some AWS services use features in other AWS services. For example, when you make a call in a service, it's common for that service to run applications in Amazon EC2 or store objects in Amazon S3. A service might do this using the calling principal's permissions, using a service role, or using a service-linked role.
 - **Forward access sessions (FAS)** – When you use an IAM user or role to perform actions in AWS, you are considered a principal. When you use some services, you might perform an action that then initiates another action in a different service. FAS uses the permissions of the principal calling an AWS service, combined with the requesting AWS service to make requests to downstream services. FAS requests are only made when a service receives a request that requires interactions with other AWS services or resources to complete. In this case, you must have permissions to perform both actions. For policy details when making FAS requests, see [Forward access sessions](#).
 - **Service role** – A service role is an [IAM role](#) that a service assumes to perform actions on your behalf. An IAM administrator can create, modify, and delete a service role from within IAM. For

more information, see [Creating a role to delegate permissions to an AWS service](#) in the *IAM User Guide*.

- **Service-linked role** – A service-linked role is a type of service role that is linked to an AWS service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your AWS account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.
- **Applications running on Amazon EC2** – You can use an IAM role to manage temporary credentials for applications that are running on an EC2 instance and making AWS CLI or AWS API requests. This is preferable to storing access keys within the EC2 instance. To assign an AWS role to an EC2 instance and make it available to all of its applications, you create an instance profile that is attached to the instance. An instance profile contains the role and enables programs that are running on the EC2 instance to get temporary credentials. For more information, see [Using an IAM role to grant permissions to applications running on Amazon EC2 instances](#) in the *IAM User Guide*.

To learn whether to use IAM roles or IAM users, see [When to create an IAM role \(instead of a user\)](#) in the *IAM User Guide*.

Managing access using policies

You control access in AWS by creating policies and attaching them to AWS identities or resources. A policy is an object in AWS that, when associated with an identity or resource, defines their permissions. AWS evaluates these policies when a principal (user, root user, or role session) makes a request. Permissions in the policies determine whether the request is allowed or denied. Most policies are stored in AWS as JSON documents. For more information about the structure and contents of JSON policy documents, see [Overview of JSON policies](#) in the *IAM User Guide*.

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

By default, users and roles have no permissions. To grant users permission to perform actions on the resources that they need, an IAM administrator can create IAM policies. The administrator can then add the IAM policies to roles, and users can assume the roles.

IAM policies define permissions for an action regardless of the method that you use to perform the operation. For example, suppose that you have a policy that allows the `iam:GetRole` action. A user with that policy can get role information from the AWS Management Console, the AWS CLI, or the AWS API.

Identity-based policies

Identity-based policies are JSON permissions policy documents that you can attach to an identity, such as an IAM user, group of users, or role. These policies control what actions users and roles can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see [Creating IAM policies](#) in the *IAM User Guide*.

Identity-based policies can be further categorized as *inline policies* or *managed policies*. Inline policies are embedded directly into a single user, group, or role. Managed policies are standalone policies that you can attach to multiple users, groups, and roles in your AWS account. Managed policies include AWS managed policies and customer managed policies. To learn how to choose between a managed policy or an inline policy, see [Choosing between managed policies and inline policies](#) in the *IAM User Guide*.

Resource-based policies

Resource-based policies are JSON policy documents that you attach to a resource. Examples of resource-based policies are IAM *role trust policies* and Amazon S3 *bucket policies*. In services that support resource-based policies, service administrators can use them to control access to a specific resource. For the resource where the policy is attached, the policy defines what actions a specified principal can perform on that resource and under what conditions. You must [specify a principal](#) in a resource-based policy. Principals can include accounts, users, roles, federated users, or AWS services.

Resource-based policies are inline policies that are located in that service. You can't use AWS managed policies from IAM in a resource-based policy.

Access control lists (ACLs)

Access control lists (ACLs) control which principals (account members, users, or roles) have permissions to access a resource. ACLs are similar to resource-based policies, although they do not use the JSON policy document format.

Amazon S3, AWS WAF, and Amazon VPC are examples of services that support ACLs. To learn more about ACLs, see [Access control list \(ACL\) overview](#) in the *Amazon Simple Storage Service Developer Guide*.

Other policy types

AWS supports additional, less-common policy types. These policy types can set the maximum permissions granted to you by the more common policy types.

- **Permissions boundaries** – A permissions boundary is an advanced feature in which you set the maximum permissions that an identity-based policy can grant to an IAM entity (IAM user or role). You can set a permissions boundary for an entity. The resulting permissions are the intersection of an entity's identity-based policies and its permissions boundaries. Resource-based policies that specify the user or role in the `Principal` field are not limited by the permissions boundary. An explicit deny in any of these policies overrides the allow. For more information about permissions boundaries, see [Permissions boundaries for IAM entities](#) in the *IAM User Guide*.
- **Service control policies (SCPs)** – SCPs are JSON policies that specify the maximum permissions for an organization or organizational unit (OU) in AWS Organizations. AWS Organizations is a service for grouping and centrally managing multiple AWS accounts that your business owns. If you enable all features in an organization, then you can apply service control policies (SCPs) to any or all of your accounts. The SCP limits permissions for entities in member accounts, including each AWS account root user. For more information about Organizations and SCPs, see [How SCPs work](#) in the *AWS Organizations User Guide*.
- **Session policies** – Session policies are advanced policies that you pass as a parameter when you programmatically create a temporary session for a role or federated user. The resulting session's permissions are the intersection of the user or role's identity-based policies and the session policies. Permissions can also come from a resource-based policy. An explicit deny in any of these policies overrides the allow. For more information, see [Session policies](#) in the *IAM User Guide*.

Multiple policy types

When multiple types of policies apply to a request, the resulting permissions are more complicated to understand. To learn how AWS determines whether to allow a request when multiple policy types are involved, see [Policy evaluation logic](#) in the *IAM User Guide*.

How Amazon Q Business works with IAM

Before you use IAM to manage access to Amazon Q, learn what IAM features are available to use with Amazon Q.

IAM features you can use with Amazon Q Business

IAM feature	Amazon Q support
Identity-based policies	Yes
Resource-based policies	No
Policy actions	Yes
Policy resources	Yes
Policy condition keys	Yes
ACLs	No
ABAC (tags in policies)	Yes
Temporary credentials	Yes
Principal permissions	Yes
Service roles	Yes
Service-linked roles	No

To get a high-level view of how Amazon Q and other AWS services work with most IAM features, see [AWS services that work with IAM](#) in the *IAM User Guide*.

Identity-based policies for Amazon Q

Supports identity-based policies	Yes
----------------------------------	-----

Identity-based policies are JSON permissions policy documents that you can attach to an identity, such as an IAM user, group of users, or role. These policies control what actions users and roles can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see [Creating IAM policies](#) in the *IAM User Guide*.

With IAM identity-based policies, you can specify allowed or denied actions and resources as well as the conditions under which actions are allowed or denied. You can't specify the principal in an identity-based policy because it applies to the user or role to which it is attached. To learn about all of the elements that you can use in a JSON policy, see [IAM JSON policy elements reference](#) in the *IAM User Guide*.

Identity-based policy examples for Amazon Q

To view examples of Amazon Q identity-based policies, see [Identity-based policy examples for Amazon Q Business](#).

Resource-based policies within Amazon Q

Supports resource-based policies	No
----------------------------------	----

Resource-based policies are JSON policy documents that you attach to a resource. Examples of resource-based policies are *IAM role trust policies* and *Amazon S3 bucket policies*. In services that support resource-based policies, service administrators can use them to control access to a specific resource. For the resource where the policy is attached, the policy defines what actions a specified principal can perform on that resource and under what conditions. You must [specify a principal](#) in a resource-based policy. Principals can include accounts, users, roles, federated users, or AWS services.

To enable cross-account access, you can specify an entire account or IAM entities in another account as the principal in a resource-based policy. Adding a cross-account principal to a resource-based policy is only half of establishing the trust relationship. When the principal and the resource are in different AWS accounts, an IAM administrator in the trusted account must also grant the principal entity (user or role) permission to access the resource. They grant permission by attaching an identity-based policy to the entity. However, if a resource-based policy grants access to a principal in the same account, no additional identity-based policy is required. For more information, see [How IAM roles differ from resource-based policies](#) in the *IAM User Guide*.

Policy actions for Amazon Q

Supports policy actions	Yes
-------------------------	-----

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The `Action` element of a JSON policy describes the actions that you can use to allow or deny access in a policy. Policy actions usually have the same name as the associated AWS API operation. There are some exceptions, such as *permission-only actions* that don't have a matching API operation. There are also some operations that require multiple actions in a policy. These additional actions are called *dependent actions*.

Include actions in a policy to grant permissions to perform the associated operation.

To see a list of Amazon Q actions, see [Actions Defined by Amazon Q Business](#) in the *Service Authorization Reference*.

Policy actions in Amazon Q use the following prefix before the action:

```
qbusiness
```

Policy actions in Amazon Q use the following prefix before the action: `qbusiness:`. For example, to grant someone permission to list an Amazon Q application with the [ListApplications](#) API operation, you include the `qbusiness:ListIndices` action in their policy. Policy statements must include either an `Action` or `NotAction` element. Amazon Q defines its own set of actions that describe tasks that you can perform with this service.

To specify multiple actions in a single statement, separate them with commas.

```
"Action": [  
  "qbusiness:action1",  
  "qbusiness:action2"  
]
```

To view examples of Amazon Q identity-based policies, see [Identity-based policy examples for Amazon Q Business](#).

Policy resources for Amazon Q

Supports policy resources	Yes
---------------------------	-----

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The Resource JSON policy element specifies the object or objects to which the action applies. Statements must include either a Resource or a NotResource element. As a best practice, specify a resource using its [Amazon Resource Name \(ARN\)](#). You can do this for actions that support a specific resource type, known as *resource-level permissions*.

For actions that don't support resource-level permissions, such as listing operations, use a wildcard (*) to indicate that the statement applies to all resources.

```
"Resource": "*"
```

To view examples of Amazon Q identity-based policies, see [Identity-based policy examples for Amazon Q Business](#).

Policy condition keys for Amazon Q

Supports service-specific policy condition keys	Yes
---	-----

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The Condition element (or Condition *block*) lets you specify conditions in which a statement is in effect. The Condition element is optional. You can create conditional expressions that use [condition operators](#), such as equals or less than, to match the condition in the policy with values in the request.

If you specify multiple Condition elements in a statement, or multiple keys in a single Condition element, AWS evaluates them using a logical AND operation. If you specify multiple values for a single condition key, AWS evaluates the condition using a logical OR operation. All of the conditions must be met before the statement's permissions are granted.

You can also use placeholder variables when you specify conditions. For example, you can grant an IAM user permission to access a resource only if it is tagged with their IAM user name. For more information, see [IAM policy elements: variables and tags](#) in the *IAM User Guide*.

AWS supports global condition keys and service-specific condition keys. To see all AWS global condition keys, see [AWS global condition context keys](#) in the *IAM User Guide*.

To see a list of Amazon Q condition keys, see [Condition Keys for Amazon Q Business](#) in the *Service Authorization Reference*. To learn with which actions and resources you can use a condition key, see [Actions Defined by Amazon Q Business](#).

To view examples of Amazon Q identity-based policies, see [Identity-based policy examples for Amazon Q Business](#).

ACLs in Amazon Q

Supports ACLs	No
---------------	----

Access control lists (ACLs) control which principals (account members, users, or roles) have permissions to access a resource. ACLs are similar to resource-based policies, although they do not use the JSON policy document format.

ABAC with Amazon Q

Supports ABAC (tags in policies)	Yes
----------------------------------	-----

Attribute-based access control (ABAC) is an authorization strategy that defines permissions based on attributes. In AWS, these attributes are called *tags*. You can attach tags to IAM entities (users or roles) and to many AWS resources. Tagging entities and resources is the first step of ABAC. Then you design ABAC policies to allow operations when the principal's tag matches the tag on the resource that they are trying to access.

ABAC is helpful in environments that are growing rapidly and helps with situations where policy management becomes cumbersome.

To control access based on tags, you provide tag information in the [condition element](#) of a policy using the `aws:ResourceTag/key-name`, `aws:RequestTag/key-name`, or `aws:TagKeys` condition keys.

If a service supports all three condition keys for every resource type, then the value is **Yes** for the service. If a service supports all three condition keys for only some resource types, then the value is **Partial**.

For more information about ABAC, see [What is ABAC?](#) in the *IAM User Guide*. To view a tutorial with steps for setting up ABAC, see [Use attribute-based access control \(ABAC\)](#) in the *IAM User Guide*.

The following table lists the actions, corresponding resource types, and condition keys for tag-based access control. Each action is authorized based on the tags associated with the corresponding resource type.

Action	Resource type	Condition keys
CreateApplication		aws:ResourceTag , aws:RequestTag , aws:TagKeys
CreateDataSource		aws:ResourceTag , aws:RequestTag , aws:TagKeys
CreateIndex		aws:ResourceTag , aws:RequestTag , aws:TagKeys
CreatePlugin		aws:ResourceTag , aws:RequestTag , aws:TagKeys
CreateRetriever		aws:ResourceTag , aws:RequestTag , aws:TagKeys
CreateWebExperience		aws:ResourceTag , aws:RequestTag , aws:TagKeys
ListTagsForResource	application, index, retriever, data source, web experience, plugin	

Action	Resource type	Condition keys
TagResource	application, index, retriever, data source, web experience, plugin	aws:ResourceTag , aws:RequestTag , aws:TagKeys
UntagResource	application, index, retriever, data source, web experience, plugin	aws:TagKeys

For information about tagging Amazon Q resources, see [Tagging resources](#). For an example identity-based policy that limits access to a resource based on resource tags, see [Tag-based policy examples](#). For more information about using tags to limit access to resources, see [Controlling access using tags](#) in the *IAM User Guide*.

Using temporary credentials with Amazon Q

Supports temporary credentials	Yes
--------------------------------	-----

Some AWS services don't work when you sign in using temporary credentials. For additional information, including which AWS services work with temporary credentials, see [AWS services that work with IAM](#) in the *IAM User Guide*.

You are using temporary credentials if you sign in to the AWS Management Console using any method except a user name and password. For example, when you access AWS using your company's single sign-on (SSO) link, that process automatically creates temporary credentials. You also automatically create temporary credentials when you sign in to the console as a user and then switch roles. For more information about switching roles, see [Switching to a role \(console\)](#) in the *IAM User Guide*.

You can manually create temporary credentials using the AWS CLI or AWS API. You can then use those temporary credentials to access AWS. AWS recommends that you dynamically generate temporary credentials instead of using long-term access keys. For more information, see [Temporary security credentials in IAM](#).

Cross-service principal permissions for Amazon Q

Supports forward access sessions (FAS)	Yes
--	-----

When you use an IAM user or role to perform actions in AWS, you are considered a principal. When you use some services, you might perform an action that then initiates another action in a different service. FAS uses the permissions of the principal calling an AWS service, combined with the requesting AWS service to make requests to downstream services. FAS requests are only made when a service receives a request that requires interactions with other AWS services or resources to complete. In this case, you must have permissions to perform both actions. For policy details when making FAS requests, see [Forward access sessions](#).

Service roles for Amazon Q

Supports service roles	Yes
------------------------	-----

A service role is an [IAM role](#) that a service assumes to perform actions on your behalf. An IAM administrator can create, modify, and delete a service role from within IAM. For more information, see [Creating a role to delegate permissions to an AWS service](#) in the *IAM User Guide*.

Warning

Changing the permissions for a service role might break Amazon Q functionality. Edit service roles only when Amazon Q provides guidance to do so.

Service-linked roles for Amazon Q

Supports service-linked roles	No
-------------------------------	----

A service-linked role is a type of service role that is linked to an AWS service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your AWS account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.

For details about creating or managing service-linked roles, see [AWS services that work with IAM](#). Find a service in the table that includes a Yes in the **Service-linked role** column. Choose the **Yes** link to view the service-linked role documentation for that service.

Identity-based policy examples for Amazon Q Business

By default, users and roles don't have permission to create or modify Amazon Q resources. They also can't perform tasks by using the AWS Management Console, AWS Command Line Interface (AWS CLI), or AWS API. To grant users permission to perform actions on the resources that they need, an IAM administrator can create IAM policies. The administrator can then add the IAM policies to roles, and users can assume the roles.

To learn how to create an IAM identity-based policy by using these example JSON policy documents, see [Creating IAM policies](#) in the *IAM User Guide*.

For details about actions and resource types defined by Amazon Q, including the format of the ARNs for each of the resource types, see [Actions, Resources, and Condition Keys for Amazon Q Business](#) in the *Service Authorization Reference*.

Topics

- [Policy best practices](#)
- [Using the Amazon Q console](#)
- [Allow users to view their own permissions](#)
- [Allow a user to converse with Amazon Q](#)
- [Allow an admin to manage plugins in an application](#)
- [Allow an admin to manage a specific plugin](#)
- [Tag-based policy examples](#)

Policy best practices

Identity-based policies determine whether someone can create, access, or delete Amazon Q resources in your account. These actions can incur costs for your AWS account. When you create or edit identity-based policies, follow these guidelines and recommendations:

- **Get started with AWS managed policies and move toward least-privilege permissions** – To get started granting permissions to your users and workloads, use the *AWS managed policies* that grant permissions for many common use cases. They are available in your AWS account. We

recommend that you reduce permissions further by defining AWS customer managed policies that are specific to your use cases. For more information, see [AWS managed policies](#) or [AWS managed policies for job functions](#) in the *IAM User Guide*.

- **Apply least-privilege permissions** – When you set permissions with IAM policies, grant only the permissions required to perform a task. You do this by defining the actions that can be taken on specific resources under specific conditions, also known as *least-privilege permissions*. For more information about using IAM to apply permissions, see [Policies and permissions in IAM](#) in the *IAM User Guide*.
- **Use conditions in IAM policies to further restrict access** – You can add a condition to your policies to limit access to actions and resources. For example, you can write a policy condition to specify that all requests must be sent using SSL. You can also use conditions to grant access to service actions if they are used through a specific AWS service, such as AWS CloudFormation. For more information, see [IAM JSON policy elements: Condition](#) in the *IAM User Guide*.
- **Use IAM Access Analyzer to validate your IAM policies to ensure secure and functional permissions** – IAM Access Analyzer validates new and existing policies so that the policies adhere to the IAM policy language (JSON) and IAM best practices. IAM Access Analyzer provides more than 100 policy checks and actionable recommendations to help you author secure and functional policies. For more information, see [IAM Access Analyzer policy validation](#) in the *IAM User Guide*.
- **Require multi-factor authentication (MFA)** – If you have a scenario that requires IAM users or a root user in your AWS account, turn on MFA for additional security. To require MFA when API operations are called, add MFA conditions to your policies. For more information, see [Configuring MFA-protected API access](#) in the *IAM User Guide*.

For more information about best practices in IAM, see [Security best practices in IAM](#) in the *IAM User Guide*.

Using the Amazon Q console

To access the Amazon Q Business console, you must have a minimum set of permissions. These permissions must allow you to list and view details about the Amazon Q resources in your AWS account. If you create an identity-based policy that is more restrictive than the minimum required permissions, the console won't function as intended for entities (users or roles) with that policy.

You don't need to allow minimum console permissions for users that are making calls only to the AWS CLI or the AWS API. Instead, allow access to only the actions that match the API operation that they're trying to perform.

To ensure that users and roles can still use the Amazon Q console, also attach the Amazon Q *ConsoleAccess* or *ReadOnly* AWS managed policy to the entities. For more information, see [Adding permissions to a user](#) in the *IAM User Guide*.

Allow users to view their own permissions

This example shows how you might create a policy that allows IAM users to view the inline and managed policies that are attached to their user identity. This policy includes permissions to complete this action on the console or programmatically using the AWS CLI or AWS API.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ViewOwnUserInfo",
      "Effect": "Allow",
      "Action": [
        "iam:GetUserPolicy",
        "iam:ListGroupsWithUser",
        "iam:ListAttachedUserPolicies",
        "iam:ListUserPolicies",
        "iam:GetUser"
      ],
      "Resource": ["arn:aws:iam::*:user/${aws:username}"]
    },
    {
      "Sid": "NavigateInConsole",
      "Effect": "Allow",
      "Action": [
        "iam:GetGroupPolicy",
        "iam:GetPolicyVersion",
        "iam:GetPolicy",
        "iam:ListAttachedGroupPolicies",
        "iam:ListGroupPolicies",
        "iam:ListPolicyVersions",
        "iam:ListPolicies",
        "iam:ListUsers"
      ],
      "Resource": "*"
    }
  ]
}
```

Allow a user to converse with Amazon Q

This example allows a user to start conversations with Amazon Q, view past conversations, and delete their conversation history for a specific Amazon Q application. The IAM context key `qbusiness:userId` is used to restrict permissions to a specific user.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "qbusiness:ChatSync",
        "qbusiness:ListMessages",
        "qbusiness:ListConversations",
        "qbusiness:DescribeExperience",
        "qbusiness>DeleteConversation"
      ],
      "Resource": [
        "arn:aws:qbusiness:<REGION>::<ACCOUNT>:application/<APPLICATION_ID>"
      ],
      "Condition": {
        "StringEquals": {
          "qbusiness:userId": "<USER_ID>"
        }
      }
    }
  ]
}
```

Allow an admin to manage plugins in an application

This example allows an Amazon Q admin to manage plugins in a chat application.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "qbusiness:CreatePlugin",
        "qbusiness:ListPlugins",
        "qbusiness:GetPlugin",

```



```

    "qbusiness:UpdatePlugin",
    "qbusiness>DeletePlugin"
  ],
  "Resource": [
    "arn:aws:qbusiness:<REGION>::<ACCOUNT>:application/<APPLICATION ID>"
  ]
}
}

```

Allow an admin to manage a specific plugin

This example allows an Amazon Q admin to manage a specific plugin.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "qbusiness:GetPlugin",
        "qbusiness:UpdatePlugin",
        "qbusiness>DeletePlugin"
      ],
      "Resource": [
        "arn:aws:qbusiness:<REGION>::<ACCOUNT>:application/<APPLICATION ID>",
        "arn:aws:qbusiness:<REGION>::<ACCOUNT>:application/<APPLICATION ID>/
plugin/<PLUGIN ID>"
      ]
    }
  ]
}

```

Tag-based policy examples

Tag-based policies are JSON policy documents that specify the actions that a principal can perform on tagged resources.

Example: Use a tag to access a resource

This example policy grants a user or role in your AWS account permission to use the ChatSync operation with any resource tagged with the key **department** and the value **finance**.

```

{

```

```

"Version": "2012-10-17",
"Statement": [
  {
    "Effect": "Allow",
    "Action": [
      "qbusiness:ChatSync"
    ],
    "Resource": [ "*" ],
    "Condition": {
      "StringEquals": {
        "aws:ResourceTag/department": "finance"
      }
    }
  }
]
}

```

Example: Use a tag to activate operations

This example policy grants a user or role in your AWS account permission to use any Amazon Q operation except the TagResource operation with any resource tagged with the key **department** and the value **finance**.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "qbusiness:*",
      "Resource": "*"
    },
    {
      "Effect": "Deny",
      "Action": [
        "qbusiness:TagResource"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:ResourceTag/department": "finance"
        }
      }
    }
  ]
}

```

```
]
}
```

Example: Use a tag to restrict access to an operation

This example policy restricts access for a user or role in your AWS account to use the ChatSync operation unless the user provides the **department** tag and it has the allowed values **finance** and **IT**.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Action": "qbusiness:ChatSync",
    "Resource": ""
  },
  {
    "Effect": "Deny",
    "Action": "qbusiness:ChatSync",
    "Resource": "",
    "Condition": {
      "Null": {
        "aws:ResourceTag/department": "true"
      }
    }
  },
  {
    "Effect": "Deny",
    "Action": "qbusiness:ChatSync",
    "Resource": "*",
    "Condition": {
      "ForAnyValue:StringNotEquals": {
        "aws:ResourceTag/department": [
          "finance",
          "IT"
        ]
      }
    }
  }
]
}
```

Troubleshooting Amazon Q Business identity and access

Use the following information to help you diagnose and fix common issues that you might encounter when working with Amazon Q and IAM.

Topics

- [I am not authorized to perform an action in Amazon Q](#)
- [I am not authorized to perform iam:PassRole](#)
- [I want to allow people outside of my AWS account to access my Amazon Q resources](#)

I am not authorized to perform an action in Amazon Q

If you receive an error that you're not authorized to perform an action, your policies must be updated to allow you to perform the action.

The following example error occurs when the mateojackson IAM user tries to use the console to view details about a fictional *my-example-widget* resource but doesn't have the fictional `qbusiness:GetWidget` permissions.

```
User: arn:aws:iam::123456789012:user/mateojackson is not authorized to perform:
qbusiness:GetWidget on resource: my-example-widget
```

In this case, the policy for the mateojackson user must be updated to allow access to the *my-example-widget* resource by using the `qbusiness:GetWidget` action.

If you need help, contact your AWS administrator. Your administrator is the person who provided you with your sign-in credentials.

I am not authorized to perform iam:PassRole

If you receive an error that you're not authorized to perform the `iam:PassRole` action, your policies must be updated to allow you to pass a role to Amazon Q.

Some AWS services allow you to pass an existing role to that service instead of creating a new service role or service-linked role. To do this, you must have permissions to pass the role to the service.

The following example error occurs when an IAM user named `marymajor` tries to use the console to perform an action in Amazon Q. However, the action requires the service to have permissions that are granted by a service role. Mary does not have permissions to pass the role to the service.

```
User: arn:aws:iam::123456789012:user/marymajor is not authorized to perform:
iam:PassRole
```

In this case, Mary's policies must be updated to allow her to perform the `iam:PassRole` action.

If you need help, contact your AWS administrator. Your administrator is the person who provided you with your sign-in credentials.

I want to allow people outside of my AWS account to access my Amazon Q resources

You can create a role that users in other accounts or people outside of your organization can use to access your resources. You can specify who is trusted to assume the role. For services that support resource-based policies or access control lists (ACLs), you can use those policies to grant people access to your resources.

To learn more, consult the following:

- To learn whether Amazon Q supports these features, see [How Amazon Q Business works with IAM](#).
- To learn how to provide access to your resources across AWS accounts that you own, see [Providing access to an IAM user in another AWS account that you own](#) in the *IAM User Guide*.
- To learn how to provide access to your resources to third-party AWS accounts, see [Providing access to AWS accounts owned by third parties](#) in the *IAM User Guide*.
- To learn how to provide access through identity federation, see [Providing access to externally authenticated users \(identity federation\)](#) in the *IAM User Guide*.
- To learn the difference between using roles and resource-based policies for cross-account access, see [How IAM roles differ from resource-based policies](#) in the *IAM User Guide*.

Compliance validation for Amazon Q Business

To learn whether an AWS service is within the scope of specific compliance programs, see [AWS services in Scope by Compliance Program](#) and choose the compliance program that you are interested in. For general information, see [AWS Compliance Programs](#).

You can download third-party audit reports using AWS Artifact. For more information, see [Downloading Reports in AWS Artifact](#).

Your compliance responsibility when using AWS services is determined by the sensitivity of your data, your company's compliance objectives, and applicable laws and regulations. AWS provides the following resources to help with compliance:

- [Security and Compliance Quick Start Guides](#) – These deployment guides discuss architectural considerations and provide steps for deploying baseline environments on AWS that are security and compliance focused.
- [Architecting for HIPAA Security and Compliance on Amazon Web Services](#) – This whitepaper describes how companies can use AWS to create HIPAA-eligible applications.

 **Note**

Not all AWS services are HIPAA eligible. For more information, see the [HIPAA Eligible Services Reference](#).

- [AWS Compliance Resources](#) – This collection of workbooks and guides might apply to your industry and location.
- [AWS Customer Compliance Guides](#) – Understand the shared responsibility model through the lens of compliance. The guides summarize the best practices for securing AWS services and map the guidance to security controls across multiple frameworks (including National Institute of Standards and Technology (NIST), Payment Card Industry Security Standards Council (PCI), and International Organization for Standardization (ISO)).
- [Evaluating Resources with Rules](#) in the *AWS Config Developer Guide* – The AWS Config service assesses how well your resource configurations comply with internal practices, industry guidelines, and regulations.
- [AWS Security Hub](#) – This AWS service provides a comprehensive view of your security state within AWS. Security Hub uses security controls to evaluate your AWS resources and to check your compliance against security industry standards and best practices. For a list of supported services and controls, see [Security Hub controls reference](#).
- [AWS Audit Manager](#) – This AWS service helps you continuously audit your AWS usage to simplify how you manage risk and compliance with regulations and industry standards.

Resilience in Amazon Q Business

The AWS global infrastructure is built around AWS Regions and Availability Zones. AWS Regions provide multiple physically separated and isolated Availability Zones, which are connected with low-latency, high-throughput, and highly redundant networking. With Availability Zones, you can design and operate applications and databases that automatically fail over between zones without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures.

For more information about AWS Regions and Availability Zones, see [AWS Global Infrastructure](#).

Infrastructure security in Amazon Q Business

As a managed service, Amazon Q Business is protected by the AWS global network security procedures that are described in the [Amazon Web Services: Overview of Security Processes](#) whitepaper.

You use AWS published API calls to access Amazon Q through the network. Clients must support the following:

- Transport Layer Security (TLS) 1.0 or later. We recommend TLS 1.2 or later.
- Cipher suites with perfect forward secrecy (PFS) such as DHE (Ephemeral Diffie-Hellman) or ECDHE (Elliptic Curve Ephemeral Diffie-Hellman). Most modern systems such as Java 7 and later support these modes.

Additionally, requests must be signed by using an access key ID and a secret access key that is associated with an IAM principal. Or, you can use the [AWS Security Token Service](#) (AWS STS) to generate temporary security credentials to sign requests.

Cross-service confused deputy prevention

The confused deputy problem is a security issue where an entity that doesn't have permission to perform an action can coerce a more-privileged entity to perform the action. In AWS, cross-service impersonation can result in the confused deputy problem. Cross-service impersonation can occur when one service (the *calling service*) calls another service (the *called service*). The calling service can be manipulated to use its permissions to act on another customer's resources in a way it should

not otherwise have permission to access. To prevent this, AWS provides tools that help you protect your data for all services with service principals that have been given access to resources in your account.

We recommend using the [aws:SourceArn](#) and [aws:SourceAccount](#) global condition context keys in resource policies to limit the permissions that Amazon Q Business gives another service to the resource. Use `aws:SourceArn` if you want only one resource to be associated with the cross-service access. Use `aws:SourceAccount` if you want to allow any resource in that account to be associated with the cross-service use.

The most effective way to protect against the confused deputy problem is to use the `aws:SourceArn` global condition context key with the full Amazon Resource Name (ARN) of the resource. If you don't know the full ARN of the resource or if you're specifying multiple resources, use the `aws:SourceArn` global context condition key with wildcard characters (*) for the unknown portions of the ARN. For example, `arn:aws:qbusiness:*:123456789012:*`.

If the `aws:SourceArn` value doesn't contain the account ID, such as an Amazon S3 bucket ARN, you must use both global condition context keys to limit permissions.

The value of `aws:SourceArn` must be `ResourceDescription`.

The following example shows how you can use the `aws:SourceArn` and `aws:SourceAccount` global condition context keys in Amazon Q to prevent the confused deputy problem.

```
{
  "Version": "2012-10-17",
  "Statement": {
    "Sid": "ConfusedDeputyPreventionExamplePolicy",
    "Effect": "Allow",
    "Principal": {
      "Service": "qbusiness.amazonaws.com"
    },
    "Action": "qbusiness:ActionName",
    "Resource": [
      "arn:aws:qbusiness::ResourceName/*"
    ],
    "Condition": {
      "ArnLike": {
        "aws:SourceArn": "arn:aws:qbusiness:*:123456789012:*"
      },
      "StringEquals": {
```



```
    "aws:SourceAccount": "123456789012"  
  }  
}  
}  
}
```

Configuration and vulnerability analysis in AWS Identity and Access Management

AWS handles basic security tasks like guest operating system (OS) and database patching, firewall configuration, and disaster recovery. These procedures have been reviewed and certified by the appropriate third parties. For more details, see the following resources:

- [Shared Responsibility Model](#)
- AWS: [Overview of Security Processes](#) (whitepaper)

The following resources also address configuration and vulnerability analysis in AWS Identity and Access Management (IAM):

- [Compliance validation for AWS Identity and Access Management](#)
- [Security best practices and use cases in AWS Identity and Access Management.](#)

Security best practices

Amazon Q Business provides several security features to consider as you develop and implement your own security policies. The following best practices are general guidelines and don't represent a complete security solution. Because these best practices might not be appropriate or sufficient for your environment, treat them as helpful considerations rather than prescriptions.

Apply principle of least privilege

Amazon Q provides a granular access policy for applications using IAM roles. We recommend that the roles be granted only the minimum set of privileges required by the job, such as covering your application and access to log destination. We also recommend auditing the jobs for permissions on a regular basis and upon any change to your application.

Role-based access control (RBAC) permissions

Administrators should strictly control role-based access control (RBAC) permissions for Amazon Q applications.

Monitoring Amazon Q Business

Monitoring is an important part of maintaining the reliability, availability, and performance of Amazon Q Business and your other AWS solutions. AWS provides the following monitoring tools to watch Amazon Q, report when something is wrong, and take automatic actions when appropriate:

- *AWS CloudTrail* captures API calls and related events made by or on behalf of your AWS account and delivers the log files to an Amazon S3 bucket that you specify. You can identify which users and accounts called AWS, the source IP address from which the calls were made, and when the calls occurred. For more information, see the [AWS CloudTrail User Guide](#).

Topics

- [Logging Amazon Q Business API calls using AWS CloudTrail](#)

Logging Amazon Q Business API calls using AWS CloudTrail

Amazon Q Business is integrated with AWS CloudTrail, a service that provides a record of actions taken by a user, role, or an AWS service in Amazon Q. CloudTrail captures all API calls for Amazon Q as events. The calls captured include calls from the Amazon Q console and code calls to the Amazon Q API operations. A *trail* enables CloudTrail to deliver log files to an Amazon S3 bucket. If you create a trail, you can enable continuous delivery of CloudTrail events to an Amazon S3 bucket, including events for Amazon Q. If you don't configure a trail, you can still view the most recent events in the CloudTrail console in **Event history**. Using the information collected by CloudTrail, you can determine the request that was made to Amazon Q, the IP address from which the request was made, who made the request, when it was made, and additional details.

For more information about CloudTrail, including how to configure and activate it, see the [AWS CloudTrail User Guide](#).

Amazon Q information in CloudTrail

CloudTrail is activated on your AWS account when you create the account. When activity occurs in Amazon Q, that activity is recorded in a CloudTrail event along with other AWS service events in **Event history**. You can view, search, and download recent events in your AWS account. For more information, see [Viewing events with CloudTrail Event history](#) in the *AWS CloudTrail User Guide*.

For an ongoing record of events in your AWS account, including events for Amazon Q, create a trail. A *trail* enables CloudTrail to deliver log files to an Amazon S3 bucket. By default, when you create a trail in the console, the trail applies to all AWS Regions. The trail logs events from all Regions in the AWS partition and delivers the log files to the Amazon S3 bucket that you specify. Additionally, you can configure other AWS services to further analyze and act upon the event data collected in CloudTrail logs. For more information, see the following topics:

- [Creating a trail for your AWS account](#)
- [CloudTrail supported services and integrations](#)
- [Configuring Amazon SNS notifications for CloudTrail](#)
- [Receiving CloudTrail log files from multiple Regions](#) and [Receiving CloudTrail log files from multiple accounts](#)

Control plane events in CloudTrail

CloudTrail supports logging the following Amazon Q actions documented in the [Amazon Q API Reference](#):

- [CreateApplication](#)
- [DeleteApplication](#)
- [GetApplication](#)
- [ListApplications](#)
- [UpdateApplication](#)
- [DeleteChatControlsConfiguration](#)
- [GetChatControlsConfiguration](#)
- [UpdateChatControlsConfiguration](#)
- [CreateDataSource](#)
- [DeleteDataSource](#)
- [GetDataSource](#)
- [ListDataSources](#)
- [UpdateDataSource](#)
- [CreateWebExperience](#)
- [DeleteWebExperience](#)

- [ListWebExperiences](#)
- [UpdateWebExperience](#)
- [CreateIndex](#)
- [DeleteIndex](#)
- [GetIndex](#)
- [ListIndices](#)
- [UpdateIndex](#)
- [CreatePlugin](#)
- [DeletePlugin](#)
- [GetPlugin](#)
- [ListPlugins](#)
- [UpdatePlugin](#)
- [CreateRetriever](#)
- [DeleteRetriever](#)
- [GetRetriever](#)
- [ListRetrievers](#)
- [UpdateRetriever](#)
- [ListTagsForResource](#)
- [TagResource](#)
- [UntagResource](#)

Every event or log entry contains information about who generated the request. The identity information helps you determine the following:

- Whether the request was made with root or AWS Identity and Access Management (IAM) user credentials.
- Whether the request was made with temporary security credentials for a role or federated user.
- Whether the request was made by another AWS service.

For more information, see [CloudTrail userIdentity element](#) in the *AWS CloudTrail User Guide*.

Data plane events in CloudTrail

[Data events](#) provide information about the resource operations performed on or in a resource (for example, reading or writing to an Amazon S3 object). These are also known as *data plane operations*. By default, CloudTrail doesn't log data events.

The following table shows the Amazon Q Business API operations logged to CloudTrail as *data events*. The **Data event type (console)** column shows the appropriate selection in the CloudTrail console. The **Amazon Q resource types** column shows the `resources.type` value that you would specify to log data events for the resource.

Data event type (console)	Amazon Q resource types	Supported data events
Amazon Q Business application	AWS::QBusiness::Application	<ul style="list-style-type: none"> ListDataSourceSyncJobs StartDataSourceSyncJob StopDataSourceSyncJob BatchPutDocument BatchDeleteDocument PutFeedback ChatSync DeleteConversation ListConversations ListMessages ListGroup DeleteGroup GetGroup PutGroup CreateUser DeleteUser GetUser UpdateUser ListDocuments

Data event type (console)	Amazon Q resource types	Supported data events
Amazon Q Business data resource	AWS::QBusiness::DataSource	<ul style="list-style-type: none"> • ListDataSourceSyncJobs • StartDataSourceSyncJob • StopDataSourceSyncJob
Amazon Q Business index	AWS::QBusiness::Index	<ul style="list-style-type: none"> • DeleteGroup • GetGroup • PutGroup • ListGroups • ListDocuments • BatchPutDocument • BatchDeleteDocument

You can log these API operations by configuring advanced event selectors to record data events for the Amazon Q resource types: `AWS::QBusiness::Application`, `AWS::QBusiness::DataSource`, and `AWS::QBusiness::Index`. To configure advanced event selectors, you can use either the CloudTrail console or the AWS CLI:

- From the CloudTrail console, choose the **Data event type** for which you want to log data events. Additionally, you can filter on the `eventName` and `resources.ARN` fields by choosing a custom log selector template. For more information, see [Logging data events with the AWS Management Console](#) in the *AWS CloudTrail User Guide*.
- From the AWS CLI, specify the `resources.type` value for which you want to log data events and set the `eventCategory` equal to `Data`. For more information, see [Logging data events with the AWS CLI](#) in the *AWS CloudTrail User Guide*.

The following example shows how to configure a trail to log all Amazon Q data events for all Amazon Q resource types.

```
aws cloudtrail put-event-selectors --trail-name trailName \
--advanced-event-selectors \
'[
  {
    "Name": "Log all data events on an Amazon Q application",
    "FieldSelectors": [
```

```

    { "Field": "eventCategory", "Equals": ["Data"] },
    { "Field": "resources.type", "Equals": ["AWS::QBusiness::Application"] }
  ]
},
{
  "Name": "Log all data events on an Amazon Q data source",
  "FieldSelectors": [
    { "Field": "eventCategory", "Equals": ["Data"] },
    { "Field": "resources.type", "Equals": ["AWS::QBusiness::DataSource"] }
  ]
},
{
  "Name": "Log all data events on an Amazon Q index",
  "FieldSelectors": [
    { "Field": "eventCategory", "Equals": ["Data"] },
    { "Field": "resources.type", "Equals": ["AWS::QBusiness::Index"] }
  ]
}
]'

```

You can additionally filter on the `eventName` and `resources.ARN` fields. For more information about configuring these fields, see [AdvancedFieldSelector](#) in the *AWS CloudTrail API Reference*.

Additional charges apply for data events. For more information about CloudTrail pricing, see [AWS CloudTrail Pricing](#).

Amazon Q management events in CloudTrail

[Management events](#) provide information about management operations that are performed on resources in your AWS account. These management events are also known as *control plane operations*. CloudTrail logs management event API operations by default.

Amazon Q Business logs the remainder of Amazon Q API operations as management events. For a list of the Amazon Q Business API operations that Amazon Q logs to CloudTrail, see the [Amazon Q API Reference](#).

Understanding Amazon Q log file entries

A trail is a configuration that enables delivery of events as log files to an Amazon S3 bucket that you specify. CloudTrail log files contain one or more log entries. An event represents a single

request from any source and includes information about the requested action, the date and time of the action, request parameters, and so on. CloudTrail log files aren't an ordered stack trace of the public API calls, so they don't appear in any specific order.

The following example shows a CloudTrail log entry that demonstrates the `CreateApplication` action.

```
{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "AssumedRole",
    "principalId": "principal ID",
    "arn": "ARN",
    "accountId": "account ID",
    "accessKeyId": "access key ID",
    "sessionContext": {
      "sessionIssuer": {
        "type": "Role",
        "principalId": "principal ID",
        "arn": "ARN",
        "accountId": "account ID",
        "userName": "user name"
      },
      "webIdFederationData": {},
      "attributes": {
        "creationDate": "yyyy-mm-ddThh:mm:ssZ",
        "mfaAuthenticated": "false"
      }
    }
  },
  "eventTime": "yyyy-mm-ddThh:mm:ssZ",
  "eventSource": "qbusiness.amazonaws.com",
  "eventName": "CreateApplication",
  "awsRegion": "region",
  "sourceIPAddress": "region",
  "userAgent": "user agent",
  "requestParameters": {
    "name": "name",
    "roleArn": "description",
    "clientToken": "client token"
  },
  "responseElements": {
    "applicationId": "application ID"
  }
}
```

```
  },
  "requestID": "request ID",
  "eventID": "event ID",
  "readOnly": false,
  "eventType": "AwsApiCall",
  "managementEvent": true,
  "recipientAccountId": "account ID",
  "eventCategory": "Management",
  "tlsDetails": {
    "tlsVersion": "TLS version",
    "cipherSuite": "cipher suite",
    "clientProvidedHostHeader": "qbusiness.us-west-2.api.aws"
  }
}
```

Service quotas for Amazon Q

The following are the service endpoints and service quotas for Amazon Q Business. To connect programmatically to Amazon Q, you use an endpoint. For more information, see [AWS service endpoints](#) in the *AWS General Reference*. Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account. For more information, see [AWS service quotas](#) in the *AWS General Reference*.

Supported Regions

The following table shows the AWS Regions and endpoints currently supported by Amazon Q.

Region name	Region	Endpoint	Protocol
US East (N. Virginia)	us-east-1	qbusiness.us-east-1.api.aws	HTTPS
US West (Oregon)	us-west-2	qbusiness.us-west-2.api.aws	HTTPS

For a list of AWS regions where Amazon Q is available, see [Amazon Q regions and endpoints](#) in the *Amazon Web Services General Reference*.

Quotas

Your AWS account has default quotas, formerly referred to as limits, for each AWS service. Unless otherwise noted, each quota is Region-specific. You can request increases for some quotas, and other quotas can't be increased.

To view the quotas for Amazon Q, open the [Service Quotas console](#). In the navigation pane, choose **AWS services** and select **Amazon Q**.

Some service quotas can be adjusted or increased. To see whether a quota can be adjusted, refer to the **Adjustable** column in the following table. To request a quota increase, use the [limit increase form](#).

The following table shows the quotas that are related to Amazon Q for your AWS account.

Name	Default	Adjustable	
Maximum number of applications per account	50	No	
Maximum number of data sources per application	50	No	
Maximum number of plugins per application	3	No	

API reference

For information on Amazon Q Business APIs, see the [Amazon Q Business API reference](#).

You can use the following AWS SDKs to access Amazon Q Business APIs:

- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for .NET](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby](#)

For information about the IAM access control permissions you need to use this API, see [IAM roles for Amazon Q Business](#) in the *Amazon Q Business User Guide*.

The following resources provide additional information about using the Amazon Q Business API:

- [Setting up for Amazon Q Business](#)
- [Amazon Q Business CLI Reference](#)
- [AWS General Reference](#)

Topics

- [Overview of Amazon Q API operations](#)

Overview of Amazon Q API operations

The following sections list Amazon Q Business API actions categorized according to functionality. Links are provided to console procedures and CLI code examples within this User Guide, along with links to corresponding operations in the *Amazon Q Business (for business use) API Reference*.

Topics

- [Creating an application](#)

- [Creating an index](#)
- [Creating a retriever](#)
- [Connecting data sources](#)
- [Upload documents directly](#)
- [Creating and customizing a web experience](#)
- [Chat and conversation management](#)
- [User and group management](#)
- [Amazon Q plugins](#)
- [Admin controls and guardrails](#)
- [User feedback](#)

Creating an application

All Amazon Q Business application actions are supported both on the console and using APIs.

- [CreateApplication](#) – See [Creating an Amazon Q Business application](#).
- [DeleteApplication](#) – See [Deleting an Amazon Q Business application](#).
- [GetApplication](#) – See [Getting Amazon Q Business application properties](#).
- [ListApplications](#) – See [Listing Amazon Q Business applications](#).
- [UpdateApplication](#) – See [Updating an Amazon Q Business application](#).

Creating an index

You can't create or manage an index using the AWS Management console. If you use the console, Amazon Q Business creates an index for you when you create an Amazon Q Business retriever. Tagging an index is the only action supported on the console.

- [CreateIndex](#) – Not supported on console.
- [DeleteIndex](#) – Not supported on console.
- [GetIndex](#) – Not supported on console.
- [ListIndices](#) – Not supported on console.
- [UpdateIndex](#) – Not supported on console.

Creating a retriever

Amazon Q Business supports retriever creation through both the console and the APIs.

- [CreateRetriever](#) – See [Creating an Amazon Q Business retriever](#) and [Creating an Amazon Kendra retriever](#).
- [DeleteRetriever](#) – See [Deleting an Amazon Q Business retriever](#) and [Deleting an Amazon Kendra retriever](#).
- [GetRetriever](#) – See [Getting Amazon Q Business retriever properties](#) and [Getting Amazon Kendra retriever properties](#).
- [ListRetrievers](#) – See [Listing retrievers](#) and [Getting Amazon Kendra retriever properties](#).
- [UpdateRetriever](#) – See [Updating an Amazon Q Business retriever](#) and [Updating an Amazon Kendra retriever](#).

Connecting data sources

Amazon Q Business supports data source connector configuration through both the console and the APIs.

- [CreateDataSource](#) – See [Configuring Amazon Q Business data source connectors](#).
- [DeleteDataSource](#) – See [Deleting a data source connector](#).
- [GetDataSource](#) – See [Getting data source connector properties](#).
- [ListDataSources](#) – See [Listing data source connectors](#).
- [UpdateDataSource](#) – See [Updating data source connectors](#).
- [StartDataSourceSyncJobs](#) – See [Starting data source connector sync jobs](#).
- [StopDataSourceSyncJobs](#) – See [Stopping data source connector sync jobs](#).
- [ListDataSourceSyncJobs](#) – See [Listing data source connector sync jobs](#).

Upload documents directly

Amazon Q Business supports direct document uploads into an Amazon Q Business index using both the console and the APIs.

- [BatchPutDocument](#) – See [Upload documents](#).
- [BatchDeleteDocument](#) – See [Deleting uploaded documents](#).

Creating and customizing a web experience

If you use the console to create your Amazon Q Business application, a web experience is created automatically and connected to your chosen data source.

- [CreateWebExperience](#) – See [Creating a web experience](#).
- [DeleteWebExperience](#) – See [Deleting an Amazon Q Business web experience](#).
- [GetWebExperience](#) – See [Getting Amazon Q Business web experience properties](#).
- [ListWebExperiences](#) – See [Listing Amazon Q Business web experiences](#).
- [UpdateWebExperience](#) – See [Updating an Amazon Q Business web experience](#).

Chat and conversation management

Chatting in an Amazon Q Business web experience preview and a deployed Amazon Q Business web experience uses the following API operations.

- [ChatSync](#) – [Preview an Amazon Q Business web experience](#).
- [DeleteConversation](#) – See [Conversation management](#).
- [ListConversations](#) – See [Conversation management](#).
- [ListMessages](#) – See [Using Amazon Q Business web experiences](#).

User and group management

Amazon Q Business provides APIs to manage users and groups in your Amazon Q Business. You can't configure user management using the console—Amazon Q Business automatically invokes these API operations for you when you configure your data source connector connection. You can use these APIs to implement your own user and group management solution if you create a Amazon Q Business application programmatically.

- [CreateUser](#) – [User mapping](#).
- [DeleteUser](#) – See [User mapping](#).
- [GetUser](#) – See [User mapping](#).
- [UpdateUser](#) – See [User mapping](#).
- [PutGroup](#) – [Group mapping](#).
- [DeleteGroup](#) – See [Group mapping](#).

- [GetGroup](#) – See [Group mapping](#).
- [ListGroup](#) – See [Group mapping](#).

Amazon Q plugins

Amazon Q Business supports plugin creation through both the console and the APIs.

- [CreatePlugin](#) – [Configuring plugins with Amazon Q Business](#).
- [DeletePlugin](#) – See [Deleting a plugin](#).
- [GetPlugin](#) – See [Getting plugin properties](#).
- [UpdatePlugin](#) – See [Updating a plugin](#).

Admin controls and guardrails

Amazon Q Business supports admin controls and guardrails configuration through both the console and the APIs.

- [UpdateChatControlsConfiguration](#) – See [Customizing global controls](#) and [Creating topic controls](#).
- [DeleteChatControlsConfiguration](#) – See [Deleting topic controls](#).
- [GetChatControlsConfiguration](#) – See [Getting topic control properties](#).

User feedback

Amazon Q Business captures end user feedback to chat responses to help address any technical issues. You can't configure this feature using the console.

- [PutFeedback](#) – See [Using web experience](#).

Document history

- **Latest documentation update:** April 16, 2024

The following table describes important changes in each release of Amazon Q Business.

Change	Description	Date
Amazon Q Business now integrates with IAM Identity Center	You can now use IAM Identity Center to manage user access for your Amazon Q application. For more details, see How Amazon Q works .	April 16, 2024
Amazon Q Business admin controls and guardrails update	The Amazon Q now supports new web experience chat modes, configurable using admin controls. For more details, see Admin controls and guardrails and Conversation settings .	April 16, 2024
Amazon Q (For Business Use) guide name update	The Amazon Q (For Business Use) Developer Guide is now called the Amazon Q Business User Guide.	March 29, 2024
Boosting chat results using document attributes	Amazon Q now supports boosting content used to generate chat responses using document attributes. For more information, see Boosting using document attributes .	February 14, 2024
Preview release	This is the initial preview release of the Amazon Q	November 28, 2023

(For Business Use) Developer Guide.