

Implementation Guide

Digital Evidence Archive



Digital Evidence Archive: Implementation 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

Overview	1
Features and benefits	2
Use cases	3
Concepts and definitions	3
Architecture overview	5
Architecture diagram	5
AWS Well-Architected design considerations	8
Operational excellence	8
Security	8
Reliability	9
Performance efficiency	9
Cost optimization	10
Sustainability	10
Architecture details	11
Web UI	11
Cases	11
User roles	12
Amazon S3	13
Landing Zone Accelerator on AWS	13
AWS services in this solution	13
Plan your deployment	16
Cost	16
Sample cost table	16
Security	19
IAM roles	19
Pre-signed Amazon S3 URLs	19
Secure Hash Algorithm (SHA256)	19
Encryption for at-rest and in-transit data	20
Audit logging	20
Supported AWS Regions	20
Quotas	21
Quotas for AWS services in this solution	21
AWS CloudFormation quotas	21
User roles quota	21

Deploy the solution	22
Deploy using AWS CloudFormation template	24
Prerequisites	24
AWS CloudFormation template	24
Step 1: Launch the stack	24
Step 2. TLS 1.2 for custom domain (optional)	26
Step 3. Add users	26
Deploy to a production environment	28
Prerequisites	29
Step 1: Clone the repository and install dependencies	33
Step 2: Customize your configuration	34
Step 3: Launch the solution	37
Step 4: Integrate your identity provider	39
Step 5: Post-deployment steps	49
Update the solution	50
Uninstall the solution	51
Using the AWS Management Console	51
Using AWS Command Line Interface	51
Deleting the Amazon S3 buckets	51
Deleting the Amazon DynamoDB tables	52
Monitoring the solution with Service Catalog AppRegistry	53
Activate CloudWatch Application Insights	54
Activate AWS Cost Explorer	55
Activate cost allocation tags associated with the solution	55
Use the solution	57
Case management	57
Cases dashboard	57
Create a case	58
Add case members to a case	58
Add or remove case member permissions	59
Upload digital evidence to a case	59
Download digital evidence from a case	60
Download a case's audit log	61
Download a file audit log for a case file	62
Deactivate a case	63
System administration	64

Assign case owner	64
Download system audit log	65
Mass data ingestion	65
Create a data vault	66
Create a source location	66
Create a destination location and file transfer task	67
Run the file transfer task in DEA	70
Associate case files	71
Disassociate case files	72
Developer guide	73
Source code	73
Integration guide	73
Adding Transport Layer Security (TLS) 1.2 support	73
Customization guide	73
Change domain names post-deployment	73
Assign users to DEA	74
Log in to Digital Evidence Archive	74
Optional security settings	74
API reference	76
Reference	77
Contributors	77
Notices	77
Revisions	79

Overview

Digital Evidence Archive on AWS (DEA) is a solution that enables investigative units to store and manage digital evidence through Amazon Simple Storage Service (Amazon S3) and Amazon DynamoDB. DEA provides a web user interface (UI) that investigators and other law enforcement personnel can use to create and update cases and associated digital evidence. DEA uses [Amazon S3 Intelligent-Tiering](#) to dynamically change the storage class used for digital evidence based on how often users access them, which reduces costs incurred when using DEA.

With DEA, law enforcement customers can optimize their total cost of ownership by reducing management of multiple storage repositories, reliance on physical devices such as USBs and hard drives, and operational costs associated with running a local data center.

DEA maintains file integrity, hashing, encryption, and audit logging to help customers meet requirements of the [Criminal Justice Information Services \(CJIS\) Security Policy](#). There are no additional charges or upfront commitments required to use DEA. You only pay for AWS services used in your DEA deployment, such as Amazon Simple Storage Service pricing. DEA integrates with your external identity provider, allowing agencies to use their existing single sign-on (SSO) configuration. This solution can also support non-standard AWS partitions, including the AWS GovCloud (US) Regions.

This implementation guide provides an overview of the Digital Evidence Archive on AWS solution, its reference architecture and components, considerations for planning the deployment, and configuration steps for deploying the Digital Evidence Archive solution to the Amazon Web Services (AWS) Cloud.

Use this navigation table to quickly find answers to these questions:

If you want to . . .	Read . . .
Know the cost for running this solution.	Cost
Understand the security considerations for this solution.	Security
Know how to plan for quotas for this solution.	Quotas

If you want to . . .	Read . . .
Know which AWS Regions are supported for this solution.	Supported AWS Regions
View or download the AWS CloudFormation template included in this solution to automatically deploy the infrastructure resources (the “stack”) for this solution.	AWS CloudFormation template

This guide is intended for solutions architects, business decision makers, DevOps engineers, data scientists, and cloud professionals who want to implement Digital Evidence Archive on AWS in their environment.

 Important

This open source AWS Solution is subject to additional notices. For more information on your customer responsibility, see [Notices](#).

Features and benefits

The Digital Evidence Archive on AWS solution provides the following features:

Easy-to-use web UI for law enforcement

A simple web UI for investigative units to manage their data in one place, without needing to use the AWS Management Console. No cloud knowledge is required to leverage the scale, elasticity, and automation capabilities of AWS through this solution.

Cost optimization with pay-as-you-go pricing

You only pay for the storage and compute services used within this solution. By default, Digital Evidence Archive uses [Amazon Simple Storage Service Intelligent-Tiering](#) to store data cost-efficiently.

Data integrity and compliance

Data within Digital Evidence Archive is encrypted. Comprehensive audit logs can be generated at the file, user, case, and system level. Access controls allow permissions to be granted on an as-needed basis. Files are hashed upon upload and can be validated to verify evidence has not been changed from its original form so users can maintain chain of custody.

Integration with [Service Catalog AppRegistry](#) and [AWS Systems Manager Application Manager](#)

This solution includes a Service Catalog AppRegistry resource to register the solution's CloudFormation template and its underlying resources as an [application](#) in both [Service Catalog AppRegistry](#) and [AWS Systems Manager Application Manager](#). With this integration, you can centrally manage the solution's resources.

Use cases

Digital Evidence Management

Law enforcement agencies struggle to store and manage an ever-increasing amount of digital evidence produced during operations and investigations. AWS and its partners build solutions for managing, storing, and analyzing digital evidence. The solutions are secure, cost-effective, and scalable to meet agencies' growing needs. By managing evidence in the cloud, maintaining compliance, preserving data integrity and the chain of custody, managing the lifecycle of evidence is easier than ever.

Concepts and definitions

This section describes key concepts and defines terminology specific to this solution:

Case Audit Log

A CSV file used to track changes made to cases created in Digital Evidence Archive.

Case Manager

An administrative user persona within Digital Evidence Archive.

Case Member

A non-admin user persona that classifies someone as being involved in a case.

Case Owner

A user persona that classifies someone as being responsible for a case.

Destination location

When performing a mass data ingestion, the destination location is the DEA dataset's location to which AWS DataSync transfers data.

Digital evidence

Digitally formatted evidence contained for a forensic case. Examples of this evidence type are photos, files, and other digital content.

Source location

When performing a mass data ingestion, the source location is the storage system or source from which AWS DataSync transfers data.

For a general reference of AWS terms, see the [AWS glossary](#) in the AWS General Reference.

Architecture overview

This section provides a reference implementation architecture diagram for the components deployed with this solution.

Architecture diagram

Deploying this solution with the default parameters deploys the following components in your AWS account.

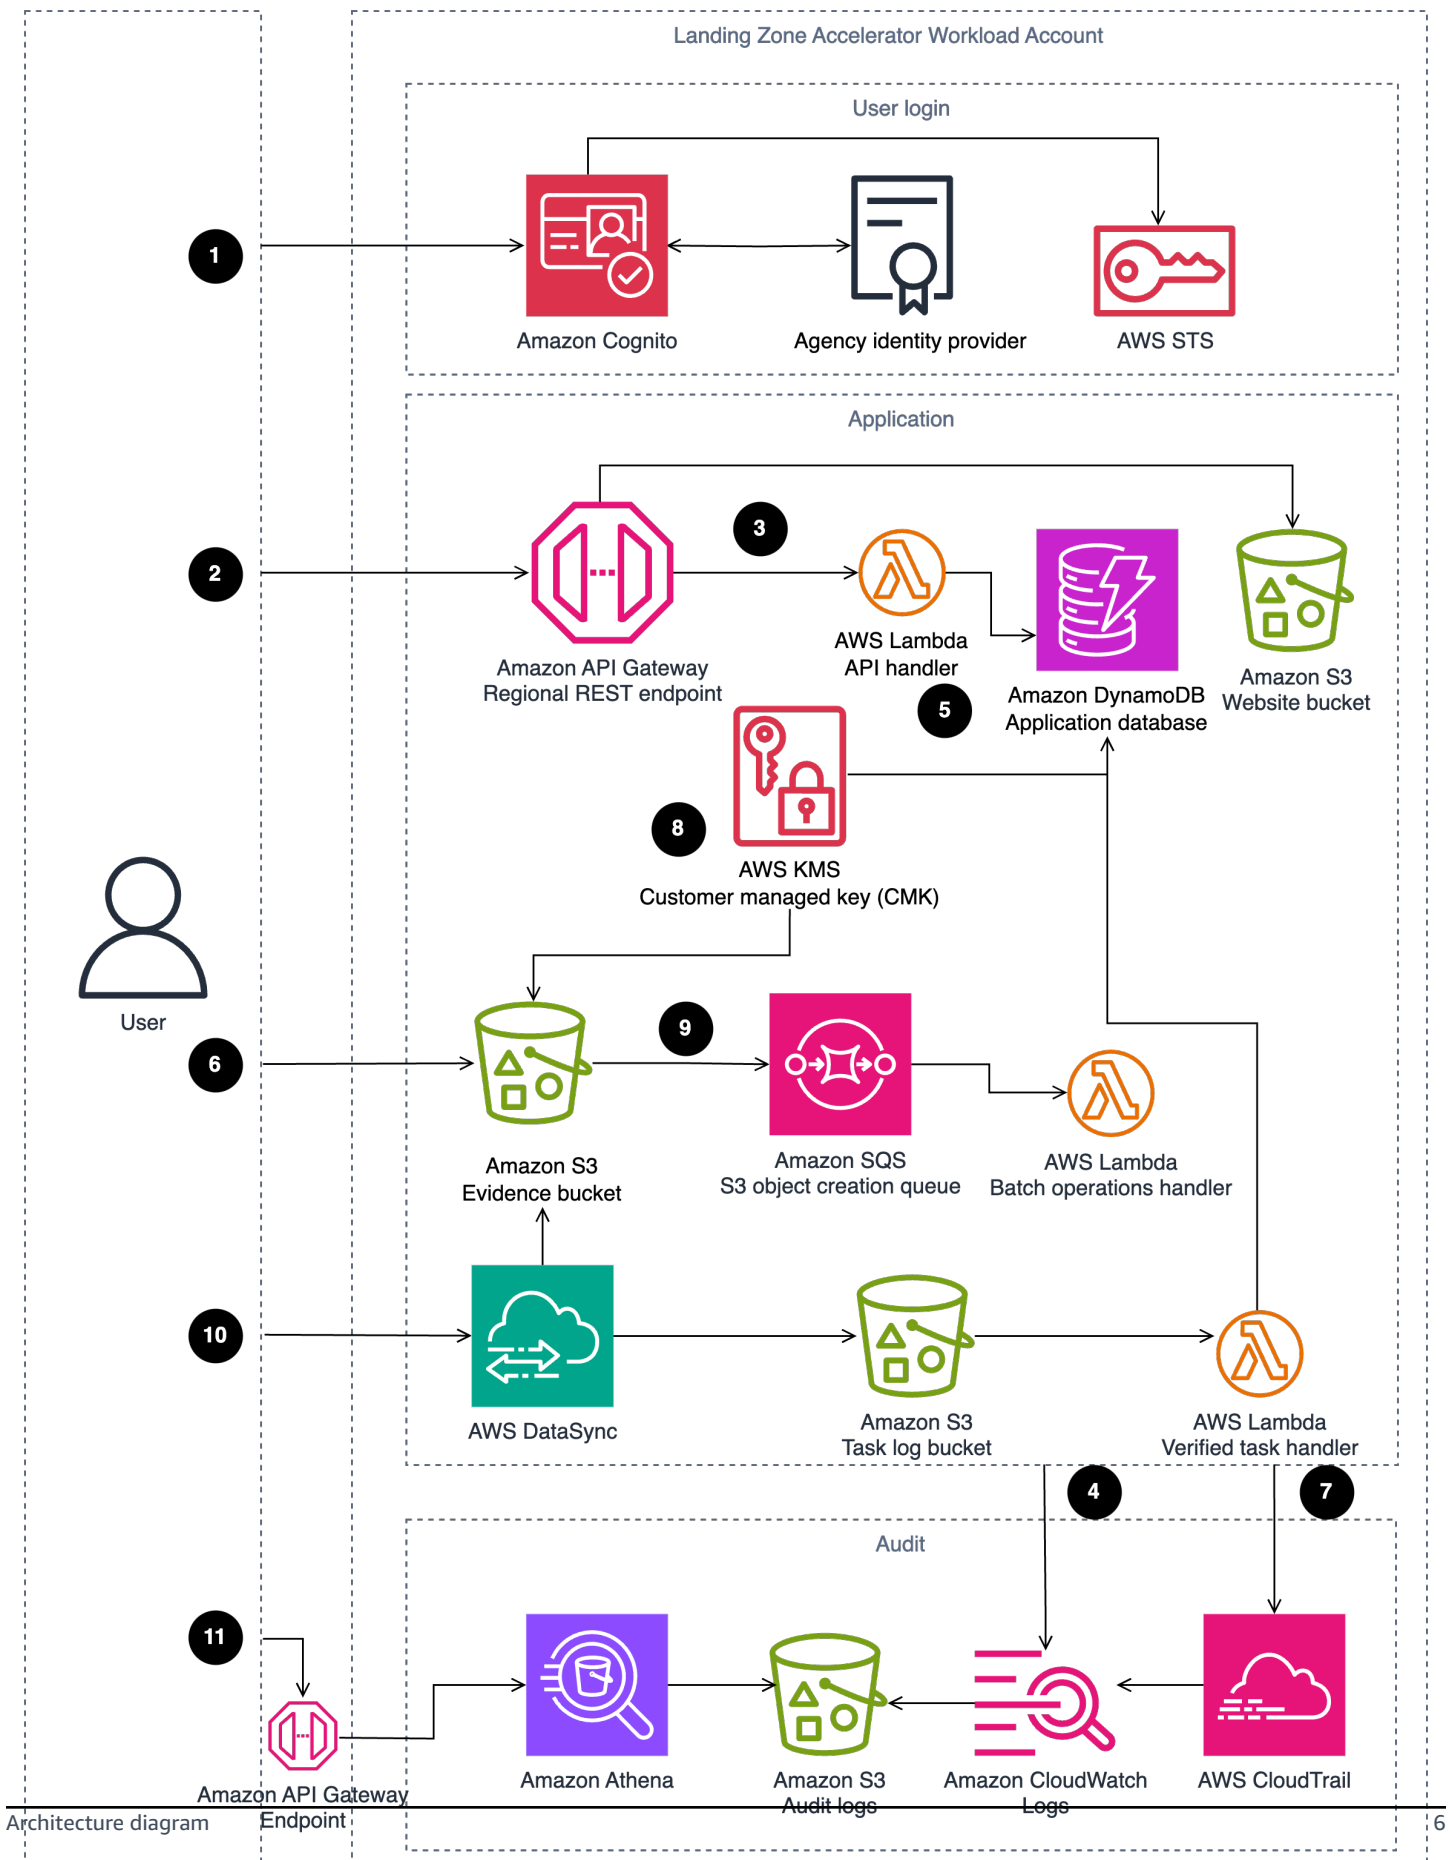


Figure 1: Digital Evidence Archive on AWS architecture

Note

AWS CloudFormation resources are created from AWS Cloud Development Kit (AWS CDK) constructs.

All communication from user to AWS goes through FIPs endpoints.

The high-level process flow for the solution components deployed with the AWS CloudFormation template into a Landing Zone Accelerator Workload Account is as follows:

1. Solution users sign in through their existing CJIS compliant Identity Provider (IdP), which federates with Amazon Cognito, to access the DEA API and web UI.
2. Users create cases using API calls to Amazon API Gateway (accessed through the web UI).
3. Case creation API calls are directed to AWS Lambda, the solution's API handler.
4. AWS Lambda sends API event data to Amazon CloudWatch for logging purposes.
5. Amazon DynamoDB registers the case creation event, and tracks user authentication sessions to mitigate malicious case actions.
6. DEA uploads data using SDK and downloads evidence using pre-signed URLs through Amazon Simple Storage Service (S3).
7. AWS CloudTrail registers CloudTrail events and Amazon S3 object-level changes in the Amazon S3 evidence bucket.
8. An AWS Key Management Service (AWS KMS) customer managed key (CMK) provides server-side encryption and prevents malicious adaptation to evidence.
9. Amazon S3 invokes AWS Lambda as needed for [S3 Batch Operations](#).
- 10 AWS DataSync receives a task to migrate data, and the reports from the migration are uploaded to the Amazon S3 tasks logs bucket. Lambda listens for the object-created event in the S3 task logs bucket and begins processing the files when detected.
- 11 Users retrieve audit reports by querying the DEA audit REST API endpoints. Amazon Athena returns case audit information to the endpoint.

AWS Well-Architected design considerations

This solution was designed with best practices from the [AWS Well-Architected Framework](#) which helps customers design and operate reliable, secure, efficient, and cost-effective workloads in the cloud.

This section describes how the design principles and best practices of the Well-Architected Framework were applied when building this solution.

Operational excellence

This section describes how the principles and best practices of the [operational excellence pillar](#) were applied when designing this solution.

- Infrastructure deployment managed through AWS CDK.
- DEA provides metrics to Amazon CloudWatch for system performance availability.
- DEA enables AWS CloudTrail for all resources, and every API call sends an event to Amazon CloudWatch. DEA provides auditing APIs at the system, user, case, and case file level that integrates both CloudTrail events (for actions taken on DEA resources outside DEA) and API events.

Security

This section describes how the principles and best practices of the [security pillar](#) were applied when designing this solution.

- Inter-service communications use AWS IAM roles.
- DEA does not pre-define user roles, but rather provides template roles as a guide. Agencies define their own roles by specifying which endpoints each role can access to best fit the agency's use case. If an endpoint is not explicitly allowed, then the user cannot call it, even if they have case permissions for that action. For example: Bob's user role does not specify the case file deletion endpoint, but one of the case's ACL specifies that he can delete files. Bob still cannot delete case files for that case.
- Amazon Cognito federates with the agency Identity Provider, and also determines which user role access to determine system access.
- Data storage, including all Amazon S3 buckets and the Amazon DynamoDB table, have encryption at rest.

- DEA uses FIPs endpoints for encryption in transit.
- Each Case has their own ACL list with multiple case actions to ensure the Principle of Least Privilege is followed.

Reliability

This section describes how the principles and best practices of the [reliability pillar](#) were applied when designing this solution.

- To ensure high availability and recovery from service failure, the solution uses AWS Serverless Services wherever possible. (Examples include AWS Lambda, Amazon API Gateway, and Amazon S3.)
- Data stored within Amazon DynamoDB has point-in-time recovery (PITR) by default to protect from accidental write or delete operations. With point-in-time recovery, you don't have to worry about creating, maintaining, or scheduling on-demand backups. Amazon DynamoDB automatically scales the database capacity based on traffic.
- Data processing uses AWS Lambda functions. AWS Lambda functions run in multiple Availability Zones to ensure that it is available to process events in case of a service interruption in a single zone by default.
- Non-testing stacks permanently retain all logging buckets by default.

Performance efficiency

This section describes how the principles and best practices of the [performance efficiency pillar](#) were applied when designing this solution.

- The solution uses AWS serverless architecture throughout. This removes the operational burden of managing physical servers, and can lower transactional costs because managed services operate at cloud scale.
- The solution can launch in any region that supports AWS services used in this solution such as: AWS Lambda, Amazon API Gateway, Amazon S3, and Amazon Cognito.
- The solution uses managed services throughout to reduce the operational burden of resource provisioning and management.

Cost optimization

This section describes how the principles and best practices of the [cost optimization pillar](#) were applied when designing this solution.

- DEA uses S3 Intelligent Tiering to lower storage costs while maintaining performance for recently accessed files.
- Because the solution uses serverless architecture, you only pay for what you use.
- Amazon DynamoDB scales capacity on demand, so you only pay for the capacity you need.

Sustainability

This section describes how the principles and best practices of the [sustainability pillar](#) were applied when designing this solution.

- To minimize the environmental impact of the backend services, DEA uses managed and serverless services. Serverless technology (such as AWS Lambda and Amazon DynamoDB) reduces carbon footprint compared to continually operating, on-premise servers.

Architecture details

This section describes the components and AWS services that make up this solution and the architecture details on how these components work together.

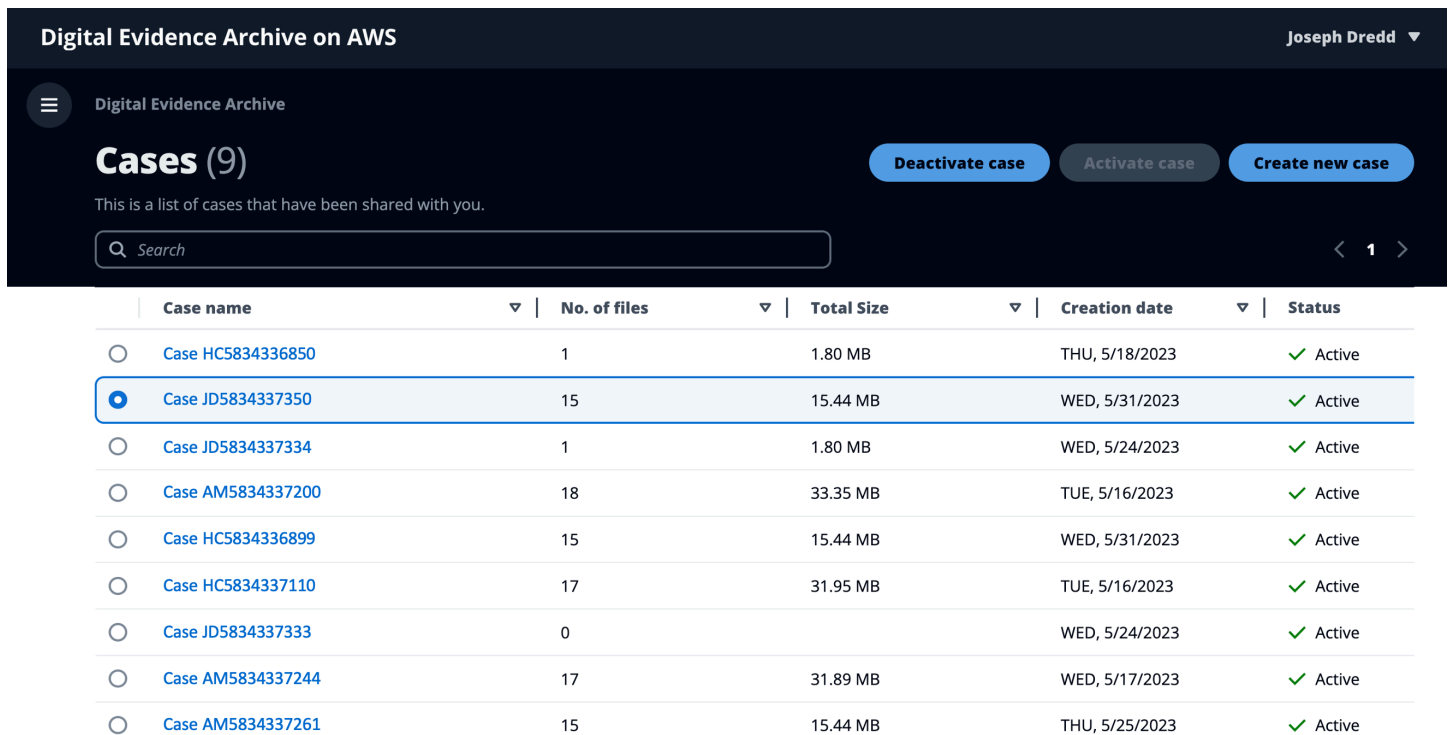
Web UI

Digital Evidence Archive on AWS (DEA) includes a pre-built [React.js](#) web UI for creating and sharing cases.

Cases

Once a case is created, the case owner or evidence manager can invite other DEA users to the case with specific case action permissions. Case actions include:

- View or update case details
- Update the case status
- Upload files
- Download files
- List case files
- Download a case audit
- Invite other users to the case
- Restore case files



The screenshot shows the 'Cases (9)' dashboard in the Digital Evidence Archive on AWS. The interface includes a search bar, navigation buttons for 'Deactivate case', 'Activate case', and 'Create new case', and a table of cases. The table has columns for Case name, No. of files, Total Size, Creation date, and Status. The case 'Case JD5834337350' is selected.

	Case name	No. of files	Total Size	Creation date	Status
<input type="radio"/>	Case HC5834336850	1	1.80 MB	THU, 5/18/2023	✓ Active
<input checked="" type="radio"/>	Case JD5834337350	15	15.44 MB	WED, 5/31/2023	✓ Active
<input type="radio"/>	Case JD5834337334	1	1.80 MB	WED, 5/24/2023	✓ Active
<input type="radio"/>	Case AM5834337200	18	33.35 MB	TUE, 5/16/2023	✓ Active
<input type="radio"/>	Case HC5834336899	15	15.44 MB	WED, 5/31/2023	✓ Active
<input type="radio"/>	Case HC5834337110	17	31.95 MB	TUE, 5/16/2023	✓ Active
<input type="radio"/>	Case JD5834337333	0		WED, 5/24/2023	✓ Active
<input type="radio"/>	Case AM5834337244	17	31.89 MB	WED, 5/17/2023	✓ Active
<input type="radio"/>	Case AM5834337261	15	15.44 MB	THU, 5/25/2023	✓ Active

Figure 1: Case dashboard

User roles

User roles are customizable. You may define the roles and allowed endpoints in the configuration file, and DEA creates them for you. For more info, see the [Developer guide](#).

Generally, there are three types of access within Digital Evidence Archive on AWS (DEA):

Evidence manager

An administrative role within DEA.

Case owner

The DEA user assigned as the owner of a particular case.

Case member/worker

DEA users assigned to a particular case. Case owners and evidence managers can add and remove members from a case, in addition to setting case permissions for members assigned to a case.

Amazon S3

Digital Evidence Archive uses Amazon S3 for cost-efficient manner storage. With S3 Intelligent-Tiering, your digital evidence is automatically stored based on how often that evidence is retrieved. Optionally, you can enable Standard, Archive, or Deep Archive access tiers when deploying or upgrading the solution, if desired.

Digital Evidence Archive uses Amazon S3 in its solution architecture for three purposes:

1. To [host and serve the web UI](#) to all solution users.
2. To store uploaded digital evidence.
3. To retain AWS CloudTrail log data on case-related changes and other CloudTrail events.

Landing Zone Accelerator on AWS

We recommend launching DEA within a Landing Zone Accelerator (LZA) workload account. Using LZA, customers with highly regulated workloads and complex compliance requirements can more easily manage and govern their multi-account environment.

AWS services in this solution

AWS Service

[Amazon Simple Storage Solution \(Amazon S3\)](#)

[Amazon DynamoDB](#)

[Amazon API Gateway](#)

Description

Core. This solution uses Amazon S3 for frontend and backend storage purposes.

Core. This solution uses a single DynamoDB table to store case and evidence metadata, in addition to access control list (ACL) definitions for cases and information on users' authentication session.

Core. This solution uses API Gateway to host public and private APIs. AWS Lambda handles these API calls to perform common solution operations.

AWS Service

[AWS Lambda](#)

Description

Core. This solution uses serverless Lambda functions, with a Node.js runtime, to handle API calls. These API events are reported to AWS CloudTrail for event logging, which includes object-level data events.

[Amazon Cognito](#)

Core. This solution uses Amazon Cognito user pools to federate with your identity provider and vend authentication tokens. This solution also uses an identity pool for authorization purposes.

[AWS Key Management Service](#)

Core. This solution uses AWS KMS keys for security and encryption purposes.

[AWS Systems Manager](#)

Supporting. This solution utilizes Systems Manager in its backend to record and analyze information on data requests that occur when the solution is deployed.

[Amazon CloudWatch](#)

Supporting. This solution is extended with CloudWatch to collect and visualize real-time logs, metrics, and event data in automated cases. Additionally, you can monitor the deployed solution's resource usage and performance issues. CloudWatch Insights is used to generate audit reports.

[AWS CloudTrail](#)

Supporting. This solution is extended with CloudTrail to monitor and record account activity while the solution is deployed. This provides you with control over storage, analysis, and remediation actions. CloudTrail can track intentional evidence modification or deletion events that occur outside of the Digital Evidence Archive solution.

AWS Service

[Amazon Data Firehose](#)

Description

Supporting. This solution uses Amazon Data Firehose to copy audit log events to Amazon S3 where they are queried for audit reporting.

[AWS Glue](#)

Supporting. This solution uses Glue to support audit data queries within Amazon S3.

[Amazon Athena](#)

Supporting. Amazon Athena extends this solution's capabilities by providing flexible querying of application audit data.

[Amazon Simple Queue Service](#)

Supporting. This solution uses Amazon SQS for asynchronous processing of application events, such as applying legal holds to newly created Amazon S3 objects.

[AWS DataSync](#)

Supporting. This solution uses DataSync for mass data ingestion.

[AWS Secrets Manager](#)

Supporting. This solution uses Secrets Manager to store secrets for authorization.

[AWS WAF](#)

Optional. This solution is extended with AWS WAF to prevent off-premises (IP) network access through the creation of an IP address allowlist that prevents unapproved IP addresses from accessing.

[AWS X-Ray](#)

Optional. This solution uses X-Ray to assist with debugging application requests.

Plan your deployment

This section describes the [cost](#), [security](#), [Regions](#), and other considerations prior to deploying the solution.

Cost

Note

You are responsible for the cost of the AWS services used while running this solution. As of March 2024, the cost for running this solution with the default settings in the US East (N. Virginia) is approximately **\$395.16 a month**.

See the pricing webpage for each AWS service used in this solution.

We recommend creating a [budget](#) through [AWS Cost Explorer](#) to help manage costs. Prices are subject to change. For full details, see the pricing webpage for each AWS service used in this solution.

Sample cost table

The following table provides a sample cost breakdown for deploying this solution with the default parameters in the US East (N. Virginia) Region for one month. In this sample, we assumed five simultaneous users in a production environment making 60 requests a minute assuming a typical 40 hour work weeks each month.

AWS Service	Dimensions	Cost [USD]
AWS CloudTrail	• 3,000,000 write management events	\$2.00 (first copy)
	• 8,000,000 read management events	\$224 (additional copies)
	• 250,000 Amazon S3 operations	
	• 250,000 DynamoDB operations	

AWS Service	Dimensions	Cost [USD]
	<ul style="list-style-type: none">• 1,500,000 Lambda operations	
AWS Key Management Service	<ul style="list-style-type: none">• 1 customer managed key• 1,535,000 symmetric requests	\$5.61
Amazon CloudWatch	<ul style="list-style-type: none">• 1 GB standard logs: data ingested• 4 metrics• 3 dashboards• 115 alarms	\$21.30
Amazon Data Firehose	<ul style="list-style-type: none">• Direct PUT• 104,857,600 records ingested per month• 5 KB per record	\$14.50
Amazon S3	<ul style="list-style-type: none">• 100 GB standard storage via PUT, COPY, POST requests with a 16 MB average object size• 59,392 GB Intelligent Tier storage with a 500 MB average object size• 4% Frequent Access Tier• 10% Infrequent Access Tier• 86% Archive Instant Access Tier	\$335.79
AWS Secrets Manager	<ul style="list-style-type: none">• 1 secret• 300 requests	\$0.40

AWS Service	Dimensions	Cost [USD]
Amazon DynamoDB	<ul style="list-style-type: none">• 1 standard table• 10 GB storage averaging 350,000 4 KB on-demand writes with 15% transactional and 84% standard• 700,000 on-demand reads with 90% eventually consistent and 10% strongly consistent	\$4.61
Amazon Athena	<ul style="list-style-type: none">• 1 query per day at 35 MB scanned per query	<\$0.01
Amazon API Gateway	<ul style="list-style-type: none">• 576,000 REST API requests with no cache	\$2.02
Amazon Simple Queue Service	<ul style="list-style-type: none">• 500,000 standard queue requests• 500,000 FIFO queue requests	No charge
AWS Lambda	<ul style="list-style-type: none">• x86 architecture• 576,000 requests• 400 ms average runtime• 2,048 MB allocated memory	\$7.80
AWS Glue	<ul style="list-style-type: none">• 1 object (table)• 300 requests	\$0.01
Amazon Cognito	<ul style="list-style-type: none">• 5 monthly active users with advanced security features enabled	\$0.25

AWS Service	Dimensions	Cost [USD]
AWS X-Ray	<ul style="list-style-type: none">576,000 requests at a 100% sampling rate	\$2.88

One-time cost to migrate bulk data:

To migrate large quantities of data using DEA's [the section called "Mass data ingestion"](#) feature and AWS DataSync, there will be a one-time charge for each migration. For example, a single **50 TB** migration costs **\$656.38**.

Security

When you build systems on AWS infrastructure, security responsibilities are shared between you and AWS. This [shared responsibility model](#) reduces your operational burden because AWS operates, manages, and controls the components including the host operating system, the virtualization layer, and the physical security of the facilities in which the services operate. For more information about AWS security, visit [AWS Cloud Security](#).

IAM roles

DEA allows agencies to customize IAM roles by modifying the configuration files for their individual use cases. To customize, specify each IAM role name for each use case, and which endpoints they should have access to, and DEA will create the roles and integrate them with the stack. See [Deploy the solution](#) for more details.

Pre-signed Amazon S3 URLs

Digital Evidence Archive uses [pre-signed Amazon S3 URLs](#) to secure digital evidence as it's downloaded. Presigned URLs protect digital transfers by providing DEA users temporary access to Amazon S3 objects for read/write purposes. Using presigned urls permits download access to a specific file location without directly vending STS credentials.

Secure Hash Algorithm (SHA256)

Digital Evidence Archive (DEA) uses secure hash algorithm 256 (SHA256) hashing to maintain data integrity for evidence stored in DEA. Hashing ensures that case data uploaded or downloaded in DEA is authentic and unaltered by malicious actors.

Encryption for at-REST and in-transit data

Case data is encrypted while at-REST and in-transit for Amazon S3 and Amazon DynamoDB. DEA is configured for [Amazon S3 server-side bucket encryption using AWS Key Management Service keys](#). In addition, Amazon S3 uses versioning and legal holds to protect against deletion.

Data in the solution's DynamoDB table also uses server-side encryption using an AWS KMS customer managed key (CMK), and through enabling [DynamoDB point-in-time recovery \(PITR\)](#). PITR provides automatic backups for data stored in the solution's deployed DynamoDB table to prevent accidental write/delete operations.

Audit logging

Digital Evidence Archive on AWS (DEA) contains built-in audit logging, so that all actions made to a case are recorded for security and assurance purposes. DEA audit logging is handled within a CSV file found within the solution's web UI. DEA audit logging is powered by AWS CloudTrail.

Supported AWS Regions

This solution uses the Amazon Cognito service, which is not currently available in all AWS Regions. You must launch this solution in an AWS Region where Amazon Cognito is available. For the most current availability of AWS services by Region, see the [AWS Regional Services List](#).

Digital Evidence Archive is supported in the following AWS Regions:

US East (Ohio)	South America (São Paulo)
US East (N. Virginia)	Europe (Paris)
US West (N. California)	Europe (Stockholm)
US West (Oregon)	Europe (Frankfurt)
Canada (Central)	Europe (Ireland) Region
Asia Pacific (Mumbai)	Europe (Milan)
Asia Pacific (Seoul)	Africa (Cape Town)
Asia Pacific (Singapore)	Middle East (Bahrain)

Asia Pacific (Sydney)	AWS GovCloud (US-East)*
Asia Pacific (Tokyo)	AWS GovCloud (US-West)

* Cognito is unavailable in AWS GovCloud (US-East). To launch DEA in that Region, you must launch the Amazon Cognito stack in AWS GovCloud (US-West) and the remaining stacks in AWS GovCloud (US-East) when deploying to a production environment. Please note, this may cause increased latency due to Cognito calls traversing different Regions.

Quotas

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account.

Quotas for AWS services in this solution

Make sure you have sufficient quota for each of the [services implemented in this solution](#). For more information, see [AWS service quotas](#).

Use the following links to go to the page for that service. To view the service quotas for all AWS services in the documentation without switching pages, view the information in the [Service endpoints and quotas](#) page in the PDF instead.

AWS CloudFormation quotas

Your AWS account has AWS CloudFormation quotas that you should be aware of when [launching the stack](#) in this solution. By understanding these quotas, you can avoid limitation errors that would prevent you from deploying this solution successfully. For more information, see [AWS CloudFormation quotas](#) in the *AWS CloudFormation User's Guide*.

User roles quota

Because Amazon Cognito user pools only allow up to 25 role mappings, you can only create 25 custom user roles.

Deploy the solution

Note

This solution uses [AWS CloudFormation templates and stacks](#) to automate its deployment. The CloudFormation template describes the AWS resources included in this solution and their properties. The CloudFormation stack provisions the resources that are described in the template.

We provide two methods of deployment:

Use the AWS CloudFormation template method for a streamlined deployment to be used for demo and testing purposes. The pre-packaged CloudFormation deployment provides a limited deployment option that contains two default roles and does not permit modification of the configuration file to define additional user roles. In addition, it does not integrate with identity providers outside of Amazon Cognito, which does not fulfill CJIS Policy Identity and Authentication requirements.

For a deployment into a production environment, we recommend using the steps in [the section called “Deploy to a production environment”](#) to define infrastructure and automate deployment. With this method, you can modify configuration files to fit your organization’s deployment and security needs.

CloudFormation template overview

Use this method for demo and testing purposes only.

Time to deploy: Approximately 15–30 minutes

[the section called “Step 1: Launch the stack”](#)

- Launch the AWS CloudFormation template into your AWS account.
- Enter values for the required parameters: CognitoDomainPrefix

[the section called “Step 2. TLS 1.2 for custom domain \(optional\)”](#)

[the section called “Step 3. Add users ”](#)

- Add users who will create and edit cases.

Production overview

Time to deploy: Approximately 30–60 minutes

[Step 1. Clone the repository](#)

- Check out the DEA repository locally
- Install necessary dependencies and build

[Step 2. Customize your configuration](#)

- Input the AWS Region and partition
- Specify your Amazon Cognito Domain
- Define your user roles
- Specify your allowed origins
- Specify whether evidence deletion is allowed
- Customize your system use notification

[Step 3. Launch the solution](#)

- Deploy DEA into your AWS account

[Step 4. Integrate your IdP](#)

- Use output from launch to integrate with your IdP
- Update your configuration file with information about your IdP
- Relaunch the stack with your updates
- Connect your user pool client to your IdP

[Step 5. Post-deployment steps](#)

Before you launch the solution, review the [cost](#), [architecture](#), [network security](#), and other considerations discussed earlier in this guide.

Deploy using AWS CloudFormation template

Note

Deployment using the AWS CloudFormation template is for demo environments only. If you plan to deploy in a production environment, see [the section called “Deploy to a production environment”](#).

Prerequisites

- You must have an AWS account.
- We additionally recommend deploying the solution within a [Landing Zone Accelerator on AWS](#) workload account.

AWS CloudFormation template

You can download the CloudFormation template for this solution before deploying it.

[View template](#)

Digital-evidence-archive - Use this template to launch the solution and all associated components. The default configuration deploys the DEA main stack and authentication, frontend, and backend resources, but you can customize the template to meet your specific needs.

Note

AWS CloudFormation resources are created from AWS Cloud Development Kit (AWS CDK) (AWS CDK) constructs.

Step 1: Launch the stack


Follow the step-by-step instructions in this section to configure and deploy the solution into your account.

Time to deploy: Approximately 30 minutes

1. 

Sign in to the AWS Management Console and select the button to launch the digital-evidence-archive AWS CloudFormation template.

2. The template launches in the US East (N. Virginia) Region by default. To launch the solution in a different AWS Region, use the Region selector in the console navigation bar.

 **Note**

This solution uses the Amazon Cognito service, which is not currently available in all AWS Regions. You must launch this solution in an AWS Region where Amazon Cognito is available. For the most current availability by Region, see the [AWS Regional Services List](#).

3. On the **Create stack** page, verify that the correct template URL is in the **Amazon S3 URL** text box and choose **Next**.
4. On the **Specify stack details** page, assign a name to your solution stack. For information about naming character limitations, see [IAM and STS Limits](#) in the *AWS Identity and Access Management User Guide*.
5. Under **Parameters**, review the parameters for this solution template and modify them as necessary. This solution uses the following default values.

Parameter	Default	Description
Amazon Cognito Domain Prefix	<i><Requires input></i>	A unique domain prefix used for your hosted Amazon Cognito login ui. Your domain name can include only lowercase letters, numbers, and hyphens. Do not use a hyphen for the first or last character. Use periods

Parameter	Default	Description
		to separate subdomain names.

- Choose **Next**.
- On the **Configure stack options** page, choose **Next**.
- On the **Review** page, review and confirm the settings. Check the box acknowledging that the template will create AWS Identity and Access Management (IAM) resources.
- Choose **Create stack** to deploy the stack.

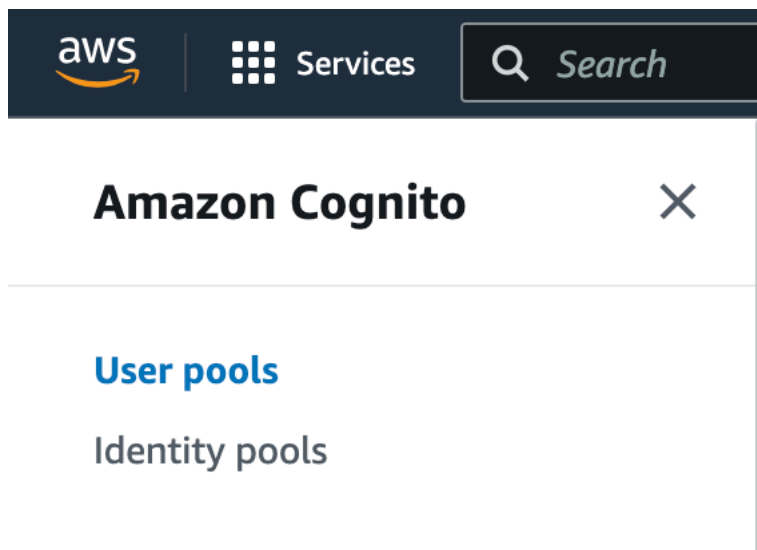
You can view the status of the stack in the AWS CloudFormation console in the **Status** column. You should receive a CREATE_COMPLETE status in approximately 30 minutes.

Step 2: TLS 1.2 for custom domain (optional)

For enhanced security, we recommend that you [add Transport Layer Security \(TLS\) 1.2 support](#) to your deployed Digital Evidence Archive solution.

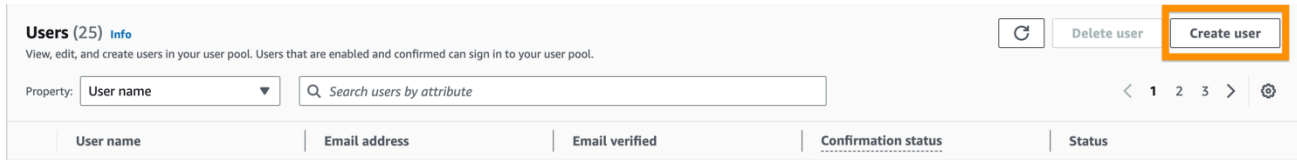
Step 3: Add users

- Open the Amazon Cognito console.
- From the navigation, choose **User pools**.



Amazon Cognito user pools

3. From the user pool list, choose the **User pool name** to open your deployment's user pool. The id can be found in your deployment's CloudFormation outputs under `DeaAuthConstructuserPoolId`.
4. From **Users**, choose **Create user**.



Users page within Amazon Cognito

5. For **User information**, enter a **User name**.
6. (Optional) To send an email invitation:
 - a. Choose **Send an email invitation**.
 - b. For **Email address**, enter the user's email address.
7. (Optional) For **Temporary password**, you may either:
 - Choose **Set a password** and enter a **Password**.
 - Choose **Generate a password**.

The user will be prompted to reset their password on their first log in.

8. Choose **Create User**.
9. From the **Users** list, choose the user you created by choosing the **User name**.
10. On the details screen, under User attributes, choose **Edit**.



User attributes screen

11. For **Required attributes**, enter **name** and **family_name**.
12. For **Optional attributes**, choose **Add attribute**.
13. For **Attribute name**, enter **custom:DEARole**.
14. For **Value**, enter either of the two following values:


- CaseWorker

Use this value for standard users who need to collaborate on cases.

- WorkingManager

Use this value for elevated users that will perform admin and CaseWorker actions.

Additional attributes

Add more user attributes. You can select from standard Cognito attributes, or assign the custom attributes that you have configured in [Sign-up experience](#) .

Attribute name

Value

custom:DEARole ▼

CaseWorker

Remove

Add another

Cancel

Save changes

Additional attributes

15. Choose **Save Changes**.

For any additional users, repeat the steps starting from step 4.

Deploy to a production environment

Use this deployment method for production environments. Alternatively, directions for production deployments may also be found in the [README.md](#) file in the GitHub repository.

Note

When deploying on your local compute environment, commands may change based on your operation system. We provide commands for both MacOS/Linux and Windows.

Topics

- [Prerequisites](#)
- [Step 1: Clone the repository and install dependencies](#)
- [Step 2: Customize your configuration](#)

- [Step 3: Launch the solution](#)
- [Step 4: Integrate your identity provider](#)
- [Step 5: Post-deployment steps](#)

Prerequisites

Important

If you are using a Windows OS to deploy, make sure that the download path for all steps does not contain spaces. This includes C:\Program Files\. Certain commands cannot run when the path contains a space.

After completing the following installation processes, restart your command prompt for the changes to take effect.

Topics

- [AWS account](#)
- [Use AWS Cloud9 for deployment \(optional\)](#)
- [Install npm and node](#)
- [Install AWS Command Line Interface](#)
- [Install Git](#)
- [Install Cygwin \(Windows only\)](#)
- [Set up a custom domain \(Optional\)](#)

AWS account

You must have 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 a user, and use only the root user to perform [tasks that require root user access](#).

Use AWS Cloud9 for deployment (optional)

We recommend using an AWS Cloud9 instance for an improved deployment experience. If you choose to use AWS Cloud9, follow these steps to configure your instance and use the Linux commands provided for your deployment.

Create the AWS Cloud9 instance

1. Open the AWS Cloud9 console at <https://console.aws.amazon.com/cloud9/>.
2. Choose **Create environment**.
3. Configure your **Environment settings**:
 - Enter a name and description for the AWS Cloud9 environment.
 - Environment type – New EC2 instance
 - Instance type – m5.large
 - Timeout – 1 hour
4. Review settings and choose **Create**.

Once the environment has been successfully created, you can open it from the AWS Cloud9 dashboard.

Modify the AWS Cloud9 volume

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. From the navigation, choose **Volumes**.
3. Select the volume associated with your AWS Cloud9 instance.
4. Choose **Modify**.
5. Increase size to at least 40 GB.
6. Choose **Modify**, and then **Modify** again to confirm.

Configure AWS Cloud9 environment

1. Verify the partition size by entering the command:

```
df -hT
```

2. Increase the partition size:

```
sudo growpart /dev/nvme0n1 1
```

3. Increase the file system inside the partition:

```
sudo xfs_growfs -d /
```

4. Verify size again:

```
df -hT
```

Clone the repo in AWS Cloud9

Before completing the remaining prerequisites, you should clone the repo:

```
git clone https://github.com/aws-solutions/digital-evidence-archive-on-aws.git
```

Install npm and node

Windows

Follow the instructions to [Install NodeJS on Windows](#). Stop before the *Install Visual Studio Code* section.

Note

During the nvm install, place npm in `C:\Users\Public\nodejs`. Additionally, we require Node 18 LTS (lts/hydrogen) instead of the latest.

MacOS/Linux

```
curl -o- https://raw.githubusercontent.com/nvm-sh/nvm/v0.35.3/install.sh | bash
```

```
source ~/.bashrc
nvm install
```

Install AWS Command Line Interface

Follow the directions to [Install or update the latest version of the AWS CLI](#) in the *AWS Command Line Interface User Guide*. Once installed, make sure you set your AWS credentials.

Important

Make sure to download to a location where the path does not have any spaces.

Install Git

Follow the directions to [Install Git](#).

Install Cygwin (Windows only)

Install [Cygwin](#) so that certain scripts can run during the build process.

Set up a custom domain (Optional)

We recommend using a custom domain for the solution in order to have a user-friendly URL. You must register a domain using Amazon Route 53 or another provider and import a certificate for the domain using AWS Certificate Manager.

Note

When accessing the DEA solution with a custom domain, you must append /ui to the URL in order to view the solution.

Route 53

1. Follow the directions to [register a domain](#) with Route 53. You should receive a confirmation email.
2. Retrieve the hosted zone for your domain. This is created automatically by Route 53.
 - a. Open the Route 53 console at <https://console.aws.amazon.com/route53/>.

- b. Choose **Hosted zones** from the navigation.
 - c. Open the hosted zone created for your domain name and copy the **Hosted zone ID**.
3. Open AWS Certificate Manager and follow these steps to [request a domain certificate](#). Verify you are in the Region where you plan to deploy the solution.
4. Choose **List certificates** from the navigation, and find your certificate request. The request should be pending.
5. Choose your **Certificate ID** to open the request.
6. From the **Domains** section, choose **Create records in Route 53**. It will take approximately ten minutes for the request to process.
7. Once the certificate is issued, copy the **ARN** from the **Certificate status** section.

Non-Route 53

1. Follow the directions for [Requesting a public certificate](#) in the same Region where you deploy the solution.

Note

To import a certificate from a third party, follow the directions for [importing certificates into AWS Certificate Manager](#).

2. Once the certificate is issued, copy the ARN to enter during [the section called “Step 2: Customize your configuration”](#).
3. Following completion of [the section called “Step 3: Launch the solution”](#), you must add a CNAME record for the domain.

Step 1: Clone the repository and install dependencies

Time: Approximately 5 minutes

Use the command line to run the following commands:

```
git clone
https://github.com/aws-solutions/digital-evidence-archive-on-aws
cd ./digital-evidence-archive-on-aws/source/
```

Install the necessary software dependencies:

```
npm install -g @microsoft/rush
npm install -g pnpm@7.16.0
npm install -g aws-cdk
```

Step 2: Customize your configuration

Time: Approximately 5 minutes

1. From the source folder in your repository, copy and rename the default configuration file.

Windows

```
cd ./common/config
copy prodexample.json prod.json
cd ../../
notepad ./common/config/prod.json
```

MacOS/Linux

```
cp ./common/config/prodexample.json ./common/config/prod.json
```

2. Inside your new configuration file, change the following fields:

- a. Specify your Region:

```
"region": "us-east-1"
```

- b. Specify a unique domain prefix for your hosted Amazon Cognito login. This is separate from your custom domain. For example:

```
"cognito": {
  "domain": "bobinohio"
},
```

- c. If you completed the prerequisites to set up a custom domain, import the `domainName`, `hostedZoneid`, `hostedZoneName`, and ACM Certificate ARN:

Route 53 domain

```
"customDomain": {
```

```
"domainName": "example.com",
"certificateArn":
"arn:aws:acm:us-east-1:ACCTNUM:certificate/CERT_NUM",
"hostedZoneId": "NJKVNFJKNVJF345903",
"hostedZoneName": "example.com"
},
```

Non-Route 53 domain

```
"customDomain": {
  "domainName": "example.com",
  "certificateArn": "arn:aws:acm:us-east-1:ACCTNUM:certificate/CERT_NUM"
},
```

d. Define your User Role Types.

You can see examples of role types already in the file and modify these endpoints or create new roles as necessary for your use case.

For each role, specify the name, description, and an array of endpoints defined by path and endpoint method. You can refer to the API Reference section of this document for a list of available endpoints. Alternatively, you can view the file called `dea-route-config.ts` under the `dea-backend` folder for the most recent list of API endpoints.

(Optional) If you intend to use [the section called “Mass data ingestion”](#), you must provide the ADMIN_ROLE_ARN with access to all Data vaults and DataSync endpoints listed in the [Digital Evidence Archive on AWS API Reference](#).

Warning

The following endpoints, which are configurable roles within the `deaRoleTypes` configuration, have elevated permissions. These endpoints permit users access to resources without having case-owner membership and are intended for admin roles:

- [Cases – List all cases](#)
- [Cases – Scoped Information](#)
- [Cases – Case owner](#)
- [Users – Audit by UserId](#)

- [Users – Audit CSV export](#)
- [System audit](#)
- [System audit – CSV export](#)

- e. If your local laws and regulations allow for or mandates the deletion of case evidence, set `deletionAllowed` field to `true`, otherwise set it to `false`.
3. Go to the front end UI to change the **System Use Notification**.

Important

The CJIS Policy Use Notification states that you must display an approved system use notification message before granting access, informing users of various usages and monitoring rules.

The message should generally discuss the following information: the user is accessing a restricted information system; system usage may be monitored, recorded, and subject to audit; unauthorized use of the system is prohibited and may be subject to criminal and/or civil penalties; use of the system indicates consent to monitoring and recording. Additionally the message shall provide appropriate privacy and security notices based on local laws and regulations. Refer to the CJIS Policy followed by your organization for the latest information.

4. To input your System Use Notification Message, open `~/digital-evidence-archive-on-aws/source/dea-ui/ui/src/common/labels.tsx` in a text editor.

Windows

```
notepad ~/digital-evidence-archive-on-aws/source/dea-ui/ui/src/common/labels.tsx
```

MacOS/Linux

```
nano ~/digital-evidence-archive-on-aws/source/dea-ui/ui/src/common/labels.tsx.
```

5. Go to the `systemUseNotificationText` definition, and change the text starting with `CUSTOMIZE YOUR SYSTEM USE NOTIFICATION TEXT...` to your approved system message.

Step 3: Launch the solution

Time: Approximately 15 minutes

1. Within the command line, open the dea-main folder.

```
cd ~/digital-evidence-archive-on-aws/source/dea-main
```

2. Export the following variables (customize as needed):

Windows

```
set STAGE=prod
set AWS_REGION=us-east-2
set DEA_CUSTOM_DOMAIN=<true if using custom domain, otherwise do NOT set>
set AWS_ACCT_NUMBER=<your 12 digit AWS account number>
```

MacOS/Linux

```
export STAGE=prod
export AWS_REGION="us-east-2"
export DEA_CUSTOM_DOMAIN=true // if using custom domain,
otherwise do not set
AWS_ACCT_NUMBER=<your 12 digit AWS account number>
```

3. (Optional) If you plan to use the [the section called "Mass data ingestion"](#) feature to import data into DEA, you must set your ADMIN_ROLE_ARN to the account on which DEA is hosted. This permits your account to see DEA's Amazon S3 datasets bucket within AWS DataSync, which is necessary for creating a location for task transfers.

Windows

```
$Env:ADMIN_ROLE_ARN=<'Your DEA AWS Account admin role arn. Example:
arn:aws:iam::<aws account number>:role/Admin'>
```

MacOS/Linux

```
export ADMIN_ROLE_ARN=<'Your DEA AWS Account admin role arn. Example:
arn:aws:iam::<aws account number>:role/Admin'>
```

4. Validate your configuration file and address any errors:

```
rushx validate:config
```

5. Install dependencies and build DEA:

```
rush cupdate  
rush build
```

6. Bootstrap the environment and deploy:

Windows

```
rushx cdk bootstrap aws://%AWS_ACCT_NUMBER%/%AWS_REGION%  
rushx cdk deploy
```

MacOS/Linux

```
rushx cdk bootstrap aws://${AWS_ACCT_NUMBER}/${AWS_REGION}  
rushx cdk deploy
```

Note

If you are using CDK deploy in the us-gov-east-1 Region, use the `--all` flag since you are deploying more than one stack. For example:

```
rushx cdk deploy --all
```

7. After the deployment completes, note the outputs for use in the next steps.
8. If you completed the optional prerequisite to [the section called "Set up a custom domain"](#) using a non-Route 53 domain, you must add a CNAME alias.
 - a. You will need the API Gateway domain name, which you can find by opening the [API Gateway console](#).
 - b. Choose **Custom domain names**.
 - c. Choose the domain name to open the details page. Under the **Configurations** tab, you will find the API Gateway domain name. It will look similar to: d-rtxxxxxxx.execute-api.us-east-1.amazonaws.com

Tip

Save your configuration file so you can reuse it when you want to update your stack. We recommend saving the configuration file in Amazon S3.

Step 4: Integrate your identity provider

Time: Approximately 10 minutes

Amazon Cognito is not CJIS compliant when used as an identity provider (IdP). You must use your CJIS compliant IdP to federate with Amazon Cognito for use in the solution. To federate, you must:

- Create an app integration in your IdP
- Relaunch the stack
- Create a custom user attribute (DEARo1e)
- Assign users to DEA using the app integration

The solution integrates with Okta, Azure Active Directory, and IAM Identity Center. With IdP integration, you can determine user access level by defining rules based on group membership or custom attributes within your IdP.

Important

IAM Identity Center does not integrate directly with DEA to support groups or custom attributes. DEA uses a PreTokenGeneration Lambda function to invoke a query to authenticate a user's membership within Identity Center. Adding groups or custom attributes directly to Identity Center will not function as expected. Follow the provided directions to integrate with Identity Center.

Integrate your IdP

Okta

Okta supports both group membership and custom attribute based authentication rules. For group membership authentication, skip the first step.

1. *(Custom attribute authentication only)* Create an attribute for users called DEARole. Limit the possible for values to only the roles configured in [the section called "Step 2: Customize your configuration"](#). For more information, see [Add custom attributes to an Okta user profile](#).
2. Add the new user pool as a SAML 2.0 enterprise application in Okta. For this, you will need your Amazon Cognito domain prefix from your configuration file and your user pool ID (listed in the CDK outputs as `DeaAuthConstructuserPoolId`). Follow the steps in [How do I set up Okta as a SAML identity provider in an Amazon Cognito user pool](#). Only complete sections *Create a SAML app in Okta* and *Configure SAML integration for your Okta app*.
 - a. For **Single sign on URL**, replace `DOMAIN_PREFIX` with your Amazon Cognito domain defined in the configuration file, and use your deployment Region for `REGION`.

Non-US Regions or Regions not using FIPs endpoints example:

```
https://DOMAIN_PREFIX.auth.REGION.amazoncognito.com/saml2/idpresponse
```

US Regions example:

```
https://DOMAIN_PREFIX.auth-fips.REGION.amazoncognito.com/saml2/idpresponse
```

- b. For the **Audience URL**, enter the Amazon Cognito URN. Replace `USER_POOL_ID` with the CDK outputs for `DeaAuthConstructuserPoolId`.

```
urn:amazon:cognito:sp:USER_POOL_ID
```

3. Configure the following attribute statements:

- firstName
- lastName
- email
- username
- deaRole

(Role created in the first step if using custom attributes.)

4. *(Group membership authentication only)* Add a group claim. For example: send all groups: Name=groups, NameFormat=Unspecified, Filter: Select Matches regex Value=.*

Azure Active Directory

Azure Active Directory supports both group membership and custom attribute based authentication rules. For group membership authentication, skip the first step.

1. *(Custom attribute authentication only)* Create an attribute for users called DEARole. Limit the possible for values to only the roles configured in [the section called “Step 2: Customize your configuration”](#). For more information, see [How to Create Custom Attributes in Azure Active Directory](#).
2. Add the new user pool as a SAML 2.0 enterprise application in Okta. For this, you will need your Amazon Cognito domain prefix from your configuration file and your user pool ID (listed in the CDK outputs as DeaAuthConstructuserPoolId). Follow the steps in [How to set up Amazon Cognito for federated authentication using Azure AD](#) to integrate Azure Active Directory. Only complete *Step 2: Add Amazon Cognito as an enterprise application in Azure AD*.
 - a. For **Single sign on URL**, replace DOMAIN_PREFIX with your Amazon Cognito domain defined in the configuration file, and use your deployment Region for REGION.

Non-US Regions or Regions not using FIPs endpoints example:

```
https://DOMAIN_PREFIX.auth.REGION.amazoncognito.com/saml2/idpresponse
```

US Regions example:

```
https://DOMAIN_PREFIX.auth-fips.REGION.amazoncognito.com/saml2/idpresponse
```

- b. For the **Audience URL**, enter the Amazon Cognito URN. Replace USER_POOL_ID with the CDK outputs for DeaAuthConstructuserPoolId.

```
urn:amazon:cognito:sp:USER_POOL_ID
```

3. Configure the following attribute statements:
 - firstName
 - lastName
 - email
 - username

- deaRole

(Role created in the first step if using custom attributes.)

4. *(Group membership authentication only)* Add a group claim. For more information on adding a group claim, see [Configure group claims for applications by using Microsoft Entra ID](#).

IAM Identity Center

Identity Center does not permit authentication of user groups or custom attributes over SAML assertions. To permit DEA to integrate with Identity Center, the solution creates a PreTokenGeneration Amazon Cognito trigger. The PreTokenGeneration trigger queries your identity store for a federated user's group memberships and adds those groups to the identity token to authorize the user.

1. Enable IAM Identity Center in your Region. You will need to choose **Enable with AWS Organizations**.
2. Connect your Active Directory or other IdP to IAM Identity Center.
 - For a self-managed directory or an AWS Managed Microsoft AD, follow the directions in [Connect to a Microsoft AD directory](#). Alternatively, you can choose to import an existing Azure Active Directory into [AWS Managed Microsoft AD](#).

Once connected, you will need to [sync your AD users in Identity Center](#).

- For an external identity provider, see <https://docs.aws.amazon.com/singlesignon/latest/userguide/manage-your-identity-source-idp.html>
3. Follow the directions to [Set up your own SAML 2.0 application](#).
 - a. Under **IAM Identity Center metadata**, copy the link for the **IAM Identity Center SAML metadata file**. You will need this when configuring SSO with DEA.
 - b. Under **Application metadata**, choose **Manually type your metadata values**.
 - c. For **Application ACS URL**, enter the Amazon Cognito domain defined in your configuration file for DOMAIN_PREFIX, and use your deployment Region for REGION.

Non-US Regions or Regions not using FIPs endpoints example:

```
https://DOMAIN_PREFIX.auth.REGION.amazoncognito.com/saml2/idpresponse
```

US Regions example:

```
https://DOMAIN_PREFIX.auth-fips.REGION.amazoncognito.com/saml2/idpresponse
```

- d. For **Application SAML audience**, enter the Amazon Cognito URN. Replace `USER_POOL_ID` with the CDK outputs for `DeaAuthConstructuserPoolId`.

```
urn:amazon:cognito:sp:USER_POOL_ID
```

4. After you create your Custom SAML 2.0 application, configure your attribute mappings.
 - a. To configure the mappings, select your application from the **Applications** list.
 - b. Choose **Actions**, and then choose **Edit attribute mappings**.
 - c. Configure the following attributes:
 - Subject → persistent → `${user:subject}`
 - firstname → basic → `${user:givenName}`
 - lastname → basic → `${user:familyName}`
 - email → basic → `${user:email}`
 - username → basic → `${user:preferredUsername}`
 - idcenterid → basic → `${user:AD_GUID}`

Configure IdP integration within DEA environment

Okta

Once you have created the SAML 2.0 integration in Okta with the appropriate User Attribute Mapping, you can start the integration process with DEA.

1. Open your configuration file in a text editor.
2. Add the following (with your specific values for each of the fields) to the configuration file.

- a. `metadataPath`

You can either link to the IdP app integration metadata via URL (recommended), or the local metadata file path.

- b. `metadataPathType`

Use either URL or FILE.

c. attributeMap

Map the Amazon Cognito field names (left) to the app integration names (right). Do not modify the Amazon Cognito field names.

i Tip

You can set a defaultRole so that if no rules map during federation, the user will be assigned the default role. If not set, the default role is no access to DEA.

Custom attribute example:

```
"idpInfo": {
  "metadataPath": <URL link to IdP metadata, or path to the file locally>
  "metadataPathType": "URL",
  "attributeMap": {
    "username": "username",
    "email": "email",
    "firstName": "firstName",
    "lastName": "lastName",
    "deaRoleName": "DEARole"
  },
  "defaultRole": 'CaseWorker'
}
```

Group membership example:

```
"idpInfo": {
  "metadataPath": "<URL link to IdP metadata>",
  "metadataPathType": "URL",
  "attributeMap": {
    "username": "username",
    "email": "email",
    "firstName": "firstname",
    "lastName": "lastname",
    "groups": "groups"
  },
  "groupToDeaRoleRules": [
```

```
{
  "filterValue": "DEAEvidenceManager",
  "deaRoleName": "EvidenceManager"
},
{
  "filterValue": "SuperUser",
  "deaRoleName": "WorkingManager"
},
{
  "filterValue": "DEA",
  "deaRoleName": "CaseWorker"
}
],
"defaultRole": 'CaseWorker'
},
```

Azure Active Directory

Once you have created the SAML 2.0 integration in Azure Active Directory with the appropriate user attribute mapping, you can now start the integration process with DEA.

1. Open your configuration file in a text editor.
2. Add the following (with your specific values for each of the fields) to the configuration file.

- a. metadataPath

You can either link to the IdP app integration metadata via URL (recommended), or the local metadata file path.

- b. metadataPathType

Use either URL or FILE.

- c. attributeMap

Map the Amazon Cognito field names (left) to the app integration names (right). Do not modify the Amazon Cognito field names.

i Tip

You can set a `defaultRole` so that if no rules map during federation, the user will be assigned the default role. If not set, the default role is no access to DEA.

Custom attribute example:

```
"idpInfo": {
  "metadataPath": <URL link to IdP metadata, or path to the file locally>
  "metadataPathType": "URL",
  "attributeMap": {
    "username": "username",
    "email": "email",
    "firstName": "firstName",
    "lastName": "lastName",
    "deaRoleName": "DEARole"
  },
  "defaultRole": 'CaseWorker'
}
```

Group membership example:

```
"idpInfo": {
  "metadataPath": "<URL link to IdP metadata>",
  "metadataPathType": "URL",
  "attributeMap": {
    "username": "username",
    "email": "email",
    "firstName": "firstname",
    "lastName": "lastname",
    "groups": "groups"
  },
  "groupToDeaRoleRules": [
    {
      "filterValue": "DEAEvidenceManager",
      "deaRoleName": "EvidenceManager"
    },
    {
      "filterValue": "SuperUser",
      "deaRoleName": "WorkingManager"
    }
  ]
}
```

```
    },  
    {  
      "filterValue": "DEA",  
      "deaRoleName": "CaseWorker"  
    }  
  ],  
  "defaultRole": 'CaseWorker'  
},
```

IAM Identity Center

Note

You will need your identity store ID and metadata path to integrate with DEA. To find your ID, go to Settings in Identity Center and choose the Identity Source tab. If you did not copy the metadata path during setup, you can find it in the Application details page.

1. Open your configuration file in a text editor.
2. Enter the metadata path and identity store ID.
3. Define your group rules. For each rule, you must define the `deaRoleName` and the `FilterValue`. The role names were defined during [the section called “Step 2: Customize your configuration”](#). The `filterValue` is a string used to assign a role to a group. For example, if a user's assigned group contains a specific string, they will be assigned to the role associated with that string.

Note

You can define up to 25 `groupToDeaRoleRules`, and they will be evaluated in order.

```
"idpInfo": {  
  "identityStoreId": "<Identity Store Id>",  
  "metadataPath": "<URL link to IdP metatdata>",  
  "metadataPathType": "URL",  
  "attributeMap": {  
    "idcenterid": "idcenterid",  
    "username": "username",
```

```
"email": "email",
"firstName": "firstname",
"lastName": "lastname"
},
"groupToDeaRoleRules": [
  {
    "filterValue": "DEAEvidenceManager",
    "deaRoleName": "EvidenceManager"
  },
  {
    "filterValue": "SuperUser",
    "deaRoleName": "WorkingManager"
  },
  {
    "filterValue": "DEA",
    "deaRoleName": "CaseWorker"
  }
]
},
```

Relaunch the stack

- Update the stack to use the information you provided in the configuration file to integrate your IdP with the DEA stack. Run the following commands:

Note

If you are using CDK deploy in the us-gov-east-1 Region, use the `--all` flag since you are deploying more than one stack. For example:

```
rushx cdk deploy --all
```

```
rush rebuild
rushx cdk deploy
```

 Tip

Save your configuration file so you can reuse it when you want to update your stack. We recommend saving the configuration file in Amazon S3.

Step 5: Post-deployment steps

We recommend [configuring your deployment](#) with the following additional steps:

- [Optional security settings](#) (highly recommended)
- [Change domain names post-deployment](#)
- [Assign users to DEA](#)
- [Log in to Digital Evidence Archive](#)

Update the solution

If you have previously deployed the solution, you can update Digital Evidence Archive on AWS to get the latest version of the solution's framework by:

1. From your source directory, pull the latest version of DEA from [GitHub](#).
2. Set the following environment variables with your deployment's values: DOMAIN_PREFIX, STAGE, AWS_REGION, ADMIN_ROLE_ARN, and DEA_CUSTOM_DOMAIN.

Windows

```
set DOMAIN_PREFIX=<your-domain-prefix>
set STAGE=prod
set AWS_REGION=us-east-2
set ADMIN_ROLE_ARN=arn:aws:iam::<aws-account-number>:role/Admin
set DEA_CUSTOM_DOMAIN=<true if using custom domain, otherwise do NOT set>
```

MacOS/Linux

```
export DOMAIN_PREFIX=<your-domain-prefix>
export STAGE=prod
export AWS_REGION="us-east-2"
export ADMIN_ROLE_ARN=arn:aws:iam::<aws-account-number>:role/Admin
export DEA_CUSTOM_DOMAIN=true // if using custom domain,
otherwise do not set
```

3. From the main directory, update and build DEA.

```
rush cupdate
rush rebuild
```

4. Redeploy DEA.

```
rushx cdk:deploy
```

5. Confirm that the details are correct and press Y to continue the update.

Uninstall the solution

You can uninstall the Digital Evidence Archive on AWS solution from the AWS Management Console or by using the AWS Command Line Interface. You must manually delete the Amazon Simple Storage Service (Amazon S3) buckets and Amazon DynamoDB tables created by this solution. AWS Solutions Implementations do not automatically delete S3 buckets and DynamoDB tables in case you have stored data to retain.

Using the AWS Management Console

1. Sign in to the [AWS CloudFormation console](#).
2. On the **Stacks** page, select this solution's installation stack.
3. Choose **Delete**.

Using AWS Command Line Interface

Determine whether the AWS Command Line Interface (AWS CLI) is available in your environment. For installation instructions, see [What Is the AWS Command Line Interface](#) in the *AWS CLI User Guide*. After confirming that the AWS CLI is available, run the following command.

```
$ aws cloudformation delete-stack --stack-name  
<installation-stack-name>
```

Deleting the Amazon S3 buckets

This solution is configured to retain the solution-created Amazon S3 bucket (for deploying in an opt-in Region) if you decide to delete the AWS CloudFormation stack to prevent accidental data loss. After uninstalling the solution, you can manually delete this S3 bucket if you do not need to retain the data. Follow these steps to delete the Amazon S3 bucket.

1. Sign in to the [Amazon S3 console](#).
2. Choose **Buckets** from the left navigation pane.
3. Locate the deamaystack S3 buckets.
4. Select the S3 bucket and choose **Delete**.

The datasets bucket (where any case uploads are stored) has additional restrictions on uploaded content, specifically the objects have object-lock and versioning enabled. To delete these buckets you will need to disable object lock on any objects and delete all versions in the bucket.

To delete non-datasets S3 buckets using AWS CLI, run the following command:

```
$ aws s3 rb s3://<bucket-name> --force
```

Note

The `--force` command empties the bucket of contents.

Deleting the Amazon DynamoDB tables

1. Sign in to the [DynamoDB console](#).
2. Select the table that starts with the same name as your stack, e.g DeaMainStack-DeaBackendStackDeaTable.
3. Choose **Delete table**.
4. Enter **delete**, and choose **Delete**.

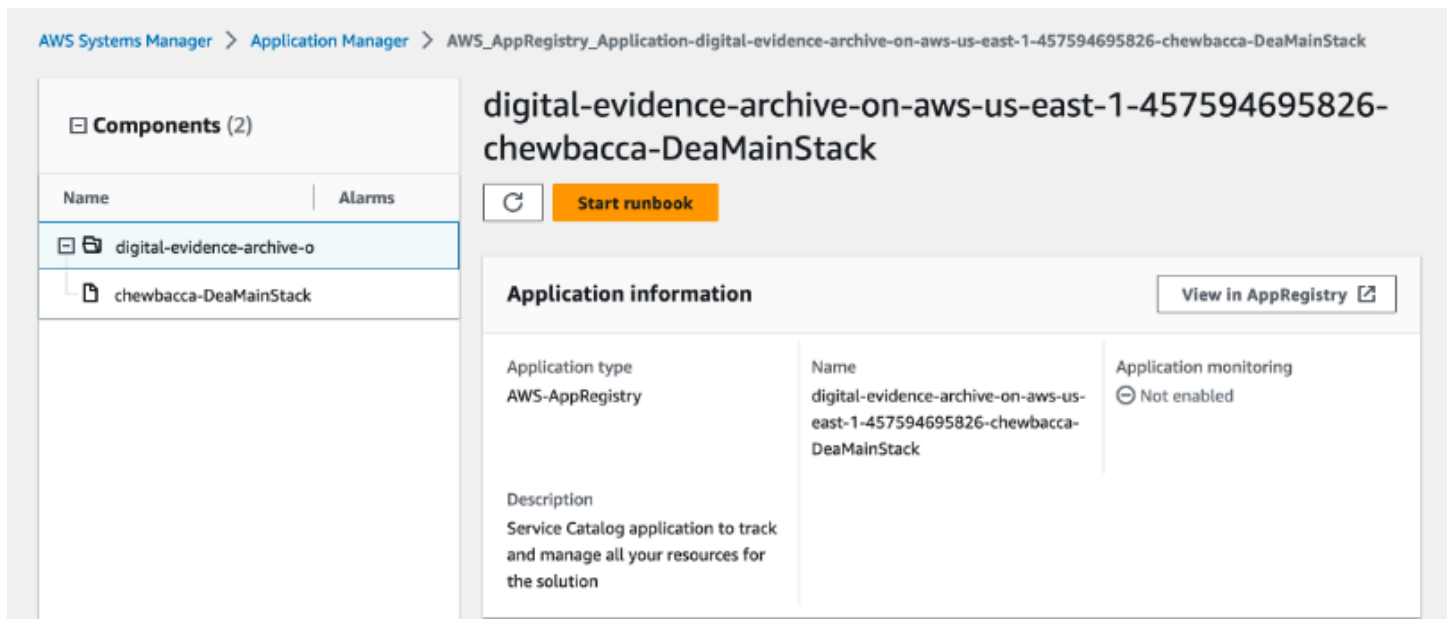
Monitoring the solution with Service Catalog AppRegistry

The Digital Evidence Archive solution includes a Service Catalog AppRegistry resource to register the CloudFormation template and underlying resources as an application in both [Service Catalog AppRegistry](#) and [AWS Systems Manager Application Manager](#).

AWS Systems Manager Application Manager gives you an application-level view into this solution and its resources so that you can:

- Monitor its resources, costs for the deployed resources across stacks and AWS accounts, and logs associated with this solution from a central location.
- View operations data for the resources of this solution in the context of an application. For example, deployment status, CloudWatch alarms, resource configurations, and operational issues.

The following figure depicts an example of the application view for the Digital Evidence Archive stack in Application Manager.



The screenshot displays the AWS Systems Manager Application Manager interface. The breadcrumb navigation at the top reads: `AWS Systems Manager > Application Manager > AWS_AppRegistry_Application-digital-evidence-archive-on-aws-us-east-1-457594695826-chewbacca-DeaMainStack`. The main heading is `digital-evidence-archive-on-aws-us-east-1-457594695826-chewbacca-DeaMainStack`. On the left, a 'Components (2)' sidebar lists `digital-evidence-archive-o` and `chewbacca-DeaMainStack`. The main content area features a 'Start runbook' button and an 'Application information' section. The 'Application information' section includes: Application type (AWS-AppRegistry), Name (digital-evidence-archive-on-aws-us-east-1-457594695826-chewbacca-DeaMainStack), Application monitoring (Not enabled), and Description (Service Catalog application to track and manage all your resources for the solution). A 'View in AppRegistry' link is also present.

Figure 1 DEA stack in Application Manager

Note

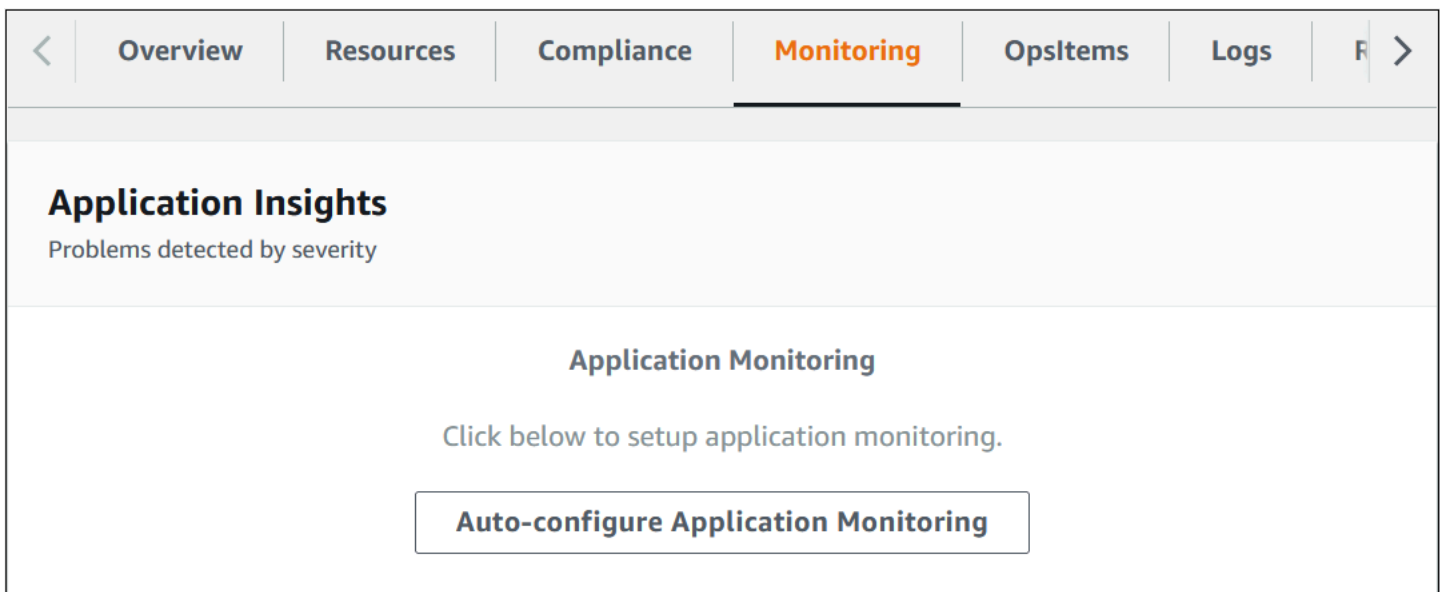
You must activate CloudWatch Application Insights, AWS Cost Explorer, and cost allocation tags associated with this solution. They are not activated by default.

Activate CloudWatch Application Insights

1. Sign in to the [Systems Manager console](#).
2. In the navigation pane, choose **Application Manager**.
3. In **Applications**, choose **AppRegistry applications**.
4. In **AppRegistry applications**, search for the application name for this solution and select it.

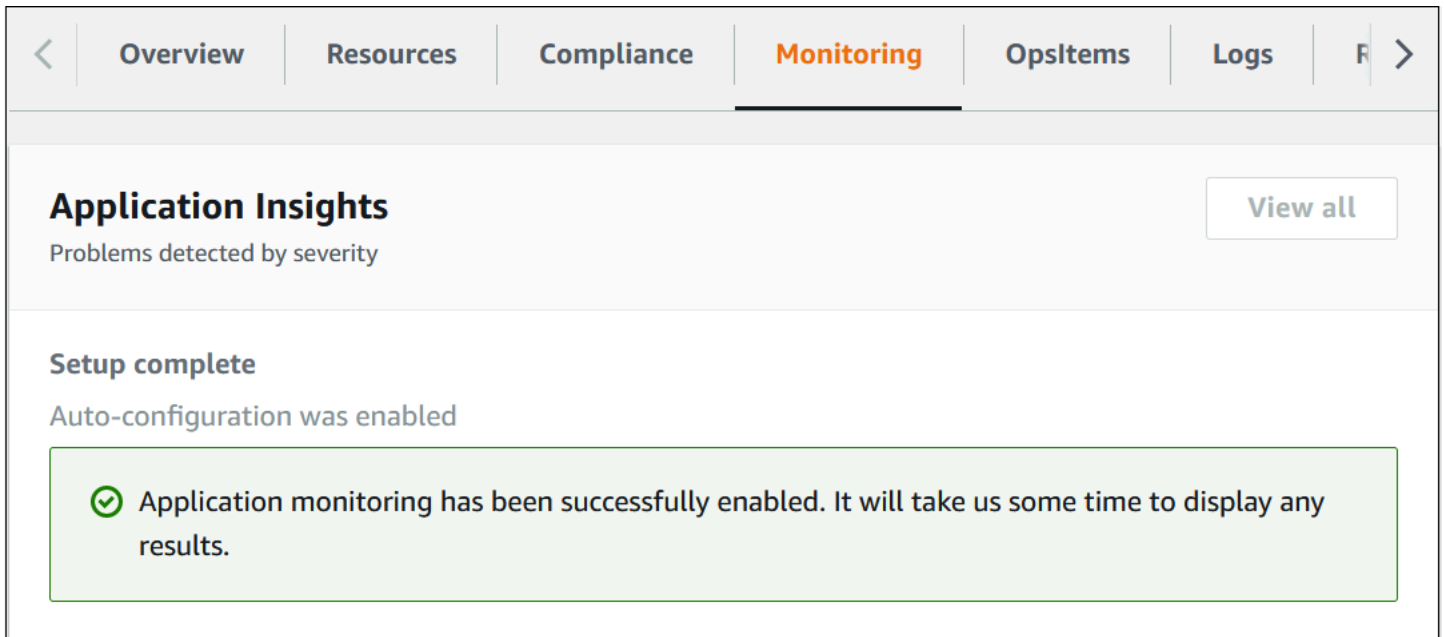
The next time you open Application Manager, you can find the new application for your solution in the **AppRegistry application** category.

5. In the **Components** tree, choose the application stack you want to activate.
6. In the **Monitoring** tab, in **Application Insights**, select **Auto-configure Application Monitoring**.



The screenshot shows the AWS Systems Manager console interface. At the top, there is a navigation bar with tabs: Overview, Resources, Compliance, Monitoring (selected), OpsItems, Logs, and a search icon. Below the navigation bar, the main content area is titled 'Application Insights' with the subtitle 'Problems detected by severity'. Underneath, there is a section for 'Application Monitoring' with the text 'Click below to setup application monitoring.' and a prominent button labeled 'Auto-configure Application Monitoring'.

Monitoring for your applications is now activated and the following status box appears:



The screenshot shows a navigation bar with tabs: Overview, Resources, Compliance, Monitoring (highlighted), OpsItems, Logs, and a search icon. Below the navigation bar, the main content area is titled "Application Insights" with a subtitle "Problems detected by severity" and a "View all" button. A "Setup complete" message states "Auto-configuration was enabled". A green success message box contains the text: "Application monitoring has been successfully enabled. It will take us some time to display any results."

Activate AWS Cost Explorer

You can see the overview of the costs associated with the application and application components within the Application Manager console through integration with AWS Cost Explorer which must be first activated. Cost Explorer helps you manage costs by providing a view of your AWS resource costs and usage over time. To activate Cost Explorer for the solution:

1. Sign in to the AWS Cost Management console.
2. In the navigation pane, select **Cost Explorer**.
3. On the **Welcome to Cost Explorer** page, choose **Launch Cost Explorer**.

The activation process can take up to 24 hours to complete. Once activated, you can open the Cost Explorer user interface to further analyze cost data for the solution.

Activate cost allocation tags associated with the solution

After you activate Cost Explorer, you must activate the cost allocation tags associated with this solution to see the costs for this solution. The cost allocation tags can only be activated from the management account for the organization. To activate cost allocation tags:

1. Sign in to the AWS Billing and Cost Management and Cost Management console.
2. In the navigation pane, select **Cost Allocation Tags**.

3. On the **Cost allocation tags** page, filter for the AppManager1CFNStackKey tag, then select the tag from the results shown.
4. Choose **Activate**.

The activation process can take up to 24 hours to complete and the tag data to appear.

Use the solution

DEA provides an intuitive user interface that helps investigative units securely manage and store digital evidence without reliance on physical devices or supporting a local data center.

In this section, you can learn about:

- [Case management](#)
- [System administration](#)
- [Mass data ingestion](#)

Case management

A Digital Evidence Archive on AWS (DEA) case consists of digital files and folders, associated case members, and permissions designated to included case members.

Cases dashboard

On login, you will be presented with a list of cases that have been shared with you.

Digital Evidence Archive on AWS

6 Digital Evidence Archive

5 Joseph Dredd

3 Deactivate case

4 Activate case

4 Create new case

Cases (9)

This is a list of cases that have been shared with you.

2 Search

	Case name	No. of files	Total Size	Creation date	Status
<input type="radio"/>	Case HC5834336850	1	1.80 MB	THU, 5/18/2023	✓ Active
<input checked="" type="radio"/>	Case JD5834337350	15	15.44 MB	WED, 5/31/2023	✓ Active
<input type="radio"/>	Case JD5834337334	1	1.80 MB	WED, 5/24/2023	✓ Active
<input type="radio"/>	Case AM5834337200	18	33.35 MB	TUE, 5/16/2023	✓ Active
<input type="radio"/>	Case HC5834336899	15	15.44 MB	WED, 5/31/2023	✓ Active
<input type="radio"/>	Case HC5834337110	17	31.95 MB	TUE, 5/16/2023	✓ Active
<input type="radio"/>	Case JD5834337333	0		WED, 5/24/2023	✓ Active
<input type="radio"/>	Case AM5834337244	17	31.89 MB	WED, 5/17/2023	✓ Active
<input type="radio"/>	Case AM5834337261	15	15.44 MB	THU, 5/25/2023	✓ Active

From the **Cases** dashboard, you can:

1. View details of a case by choosing the case name.
2. Search for a case.
3. Deactivate a case.
4. Create a new case.
5. Use the dropdown to sign out.
6. View the navigation panel.

Create a case

1. From the **Cases** dashboard, choose **Create a new case**.
2. Enter case details.
 - a. Enter a **Case name**.
 - b. **(Optional)** Enter a description.
3. Choose **Create**.

Digital Evidence Archive > Create case

Create case

All fields are required unless specified.

Enter Case Details 2

Case name
Create a Unique name that you can easily reference.

Alphanumeric characters only. No special characters.

Description - optional
Enter a brief description for your case.

3

Cancel Create

Add case members to a case

1. From the **Cases** dashboard, open a case for which you want to add members.
2. From the **Case Details** page, choose the **Assign case permissions** tab.

Note

Only members who have signed in to the solution will be available. If a member you are searching for is unavailable, they must sign in to the solution. You can try adding them again after they have signed in.

3. Search for and select users to add as case members by entering names or emails.
4. Choose **Add**.
5. Give case action permissions to individual case members from the **Permissions** dropdown.
6. Choose **Save updates**.

The screenshot shows the 'Assign case permissions' interface. At the top, there are tabs for 'Case files' and 'Assign case permissions' (highlighted with a blue underline and a yellow circle with the number 2). Below the tabs is the 'Case Members' section. It features a search bar labeled 'Search for people' (with a yellow circle 3) and a search input field with the placeholder 'Search by name or email'. To the right of the search bar is an 'Add' button (with a yellow circle 4). Below the search bar, there is a list of case members. The first member is 'UXtest Manager', and the second is 'UXtest One'. For each member, there is a 'Permission(s)' dropdown menu (with a yellow circle 5) and a 'Remove' button. The 'UXtest Manager' dropdown shows 'Choose permissions', 'View case X', and a '+8' button. The 'UXtest One' dropdown shows 'Choose permissions' and a 'Remove' button. At the bottom right of the interface is a 'Save updates' button (with a yellow circle 6).

Add or remove case member permissions

Additionally, you can add permissions to an existing user by choosing the **[+]** to add permissions.

To remove permissions from an existing user, choose the **X** next to the permission you want to remove.

Upload digital evidence to a case

1. From the **Cases** dashboard, open the case for which you want to add evidence.

2. From the **Case Files** tab, choose **Upload**.
3. Enter a description.
4. Enter a reason this evidence is being uploaded.
5. To add evidence you may:
 - Drag and drop files
 - Choose files.
 - Choose folders.
6. Choose **Upload and save**.
7. Choose **Done** once you have finishing uploading.

Uploaded evidence is stored in Amazon S3, using [S3 Object Lock](#) to prevent evidence manipulation, and associated to the case in the DEA solution's database.

Upload folders and files

Upload details

All fields are required unless otherwise indicated.

Description 3

Enter a brief description of the evidence being uploaded. Max character limit 250.

A description of the evidence is here.

Reason for uploading evidence 4

Explain why you're accessing the case files.

This evidence is important.

5

Drag and drop files or

 Choose folders

 Choose files

All file types accepted. 5TB maximum file size.

Done

6 Upload and save

Download digital evidence from a case

1. From the **Cases** dashboard, open the case from which you want to download evidence.
2. Choose the files you want to download.

- To download all files, choose the top check box.
- To download specific files, select each file you want to download.

3. Choose **Download**.

The screenshot displays the 'Case JD5834337350' page in the Digital Evidence Archive on AWS. The page is divided into two main sections: 'Case details' and 'Case files'.

Case details: This section includes an 'Edit' button and three columns of information:

- Creation date:** 5/31/2023, 1:39:37 PM
- Description:** Test
- Status:** Download Case Audit Log CSV

Case files: This section has tabs for 'Case files' (selected) and 'Case members'. It features an 'Upload' button, a 'Download' button with a notification badge '3', and a search bar labeled 'Search by file name'. Below the search bar is a table of case files:

<input type="checkbox"/>	File name ▾	File type ▾	Size ▾	Upload date ▾	Uploaded by ▾	Status ▾
<input type="checkbox"/>	Calendar Log.rtf	text/rtf	16.92 KB	2023-05-31 T 17:40...	Joseph Dredd	Active

Download a case's audit log

Every case within Digital Evidence Archive has an associated **Case Audit Log**, formatted as a CSV file. This audit log tracks all changes made to a case to provide better case visibility and prevent malicious or accidental case changes.

1. From the **Cases** dashboard, open the case for which you want to download a case audit log.
2. From the **Case Details** section, choose **Download Case Audit Log CSV**.

Case Details Edit

Creation Date 5/24/2023, 2:18:15 PM	Description Case JD 5834337333	Audit log 2 Download Case Audit Log CSV
Status ✓ Active		

Download a file audit log for a case file

If a case created in Digital Evidence Archive has case files associated to it, you can download a file audit log for selected case files.

1. From the **Cases** dashboard, open the case for which you want to download a case file audit log.
2. From the **Case Files** section, select any number of case files to add to the case file audit.
3. Choose **Case File Audit Log**.

The file audit log will be downloaded to your local device.

Case Files | Case members 3

Case Files /

Uploaded folders/files associated with this case.

Upload | Download | Case File Audit Log

Search by file name < 1 >

<input type="checkbox"/>	Name	File type	Size	Upload date	Uploaded by	Status
<input type="checkbox"/> 2	Screenshot 2023-06-08 at 3.19.10 PM.png	image/png	82.49 KB	2023-06-08T23:42...	UXtest Manager	ACTIVE
<input type="checkbox"/>	Screenshot 2023-06-08 at 3.24.45 PM.png	image/png	44.91 KB	2023-06-08T23:42...	UXtest Manager	ACTIVE
<input type="checkbox"/>	Screenshot 2023-06-08 at 4.36.23 PM.png	image/png	78.78 KB	2023-06-08T23:42...	UXtest Manager	ACTIVE
<input type="checkbox"/>	Screenshot					

Deactivate a case

If a case must be locked down for access, or the case is closed, you can deactivate the case. Additionally, if you specified that deletion is allowed for evidence, you can use deactivate to delete all case files.

1. From the **Cases** dashboard, select the case you want to deactivate.
2. Choose **Deactivate Case**.

Digital Evidence Archive on AWS Danielle Melvin ▾

Digital Evidence Archive

Cases (1)

This is a list of cases that have been shared with you.

Search

Case name	Number of files	Total Size	Creation date
1 <input checked="" type="radio"/> The Missing Checks	2	461.68 KB	TUE, 6/27/2023

3. If you want to delete all files, choose **Delete all files**.
4. Choose **Deactivate**.

Digital Evidence Archive on AWS Drew Pecka ▾

Digital Evidence Archive

My cases (3)

Search for cases, view case details, or create new cases to store digital evidence.

Search

Case name	Number of files	Total size	Creation date	Status
<input type="radio"/> The One Case	0	-	July 9, 2023	Inactive
<input type="radio"/> Another Case				Inactive
<input checked="" type="radio"/> Case to be Deactivated				Active

Are you sure you want to deactivate Case to be Deactivated?

Once the case is deactivated, anyone with access will not be able to edit, add case members, or upload and download files to the case. We will keep all your files unless you prefer to delete them.

Delete all files **3**

4


Cancel **Deactivate**

You can still view the case in the dashboard, and if you open the **Case Details**, you can still:

- See a list of files.

- Download a case audit log.

You will no longer be able to upload or download files within the case.

 **Note**

All deleted files will display as deleted. Folders will continue to function as active while displaying as deleted in order to see a list of files.

System administration

For tasks requiring elevated permission, you must use the Working Manager role created during deployment.


This user can:

- [Assign case owner](#)
- [Download system audit log](#)

Assign case owner

If a case owner is re-assigned or leaves the agency, you must re-assign the case to another owner.

1. From the navigation, choose **All cases**.
2. Choose the case for which you plan to reassign ownership.
3. Search for and select the name of the new owner.


 **Note**

Only members who have signed in to the solution will be available. If a member you are searching for is unavailable, they must sign in to the solution. You can try adding them again after they have signed in.

4. Choose **Add**.

Download system audit log

From the navigation, choose **Download System Audit Log**.



Mass data ingestion

Important

To perform a mass data ingestion, you must have the necessary permissions provided in the ADMIN_ROLE_ARN. If you are using the default deployment permissions, this role is called EvidenceManager. For custom roles, it must be a role with full access to the data vault APIs. For more information on setting the role, see [the section called "Step 3: Launch the solution"](#).

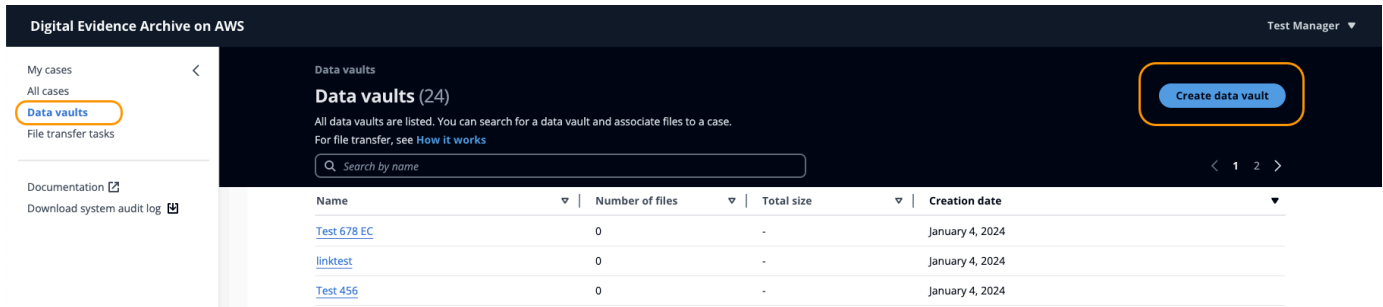
Digital Evidence Archive provides mass data ingestion to streamline the ingestion of large quantities of data into the solution using AWS DataSync. With this feature, you can create a data vault and associate it with external data archives to maintain data ingestion over time through task creation. Once data is imported, you can associate evidence to cases. Though data may be uploaded directly through the user interface, this is not a recommended practice for bulk data migrations.

To perform a mass data ingestion:

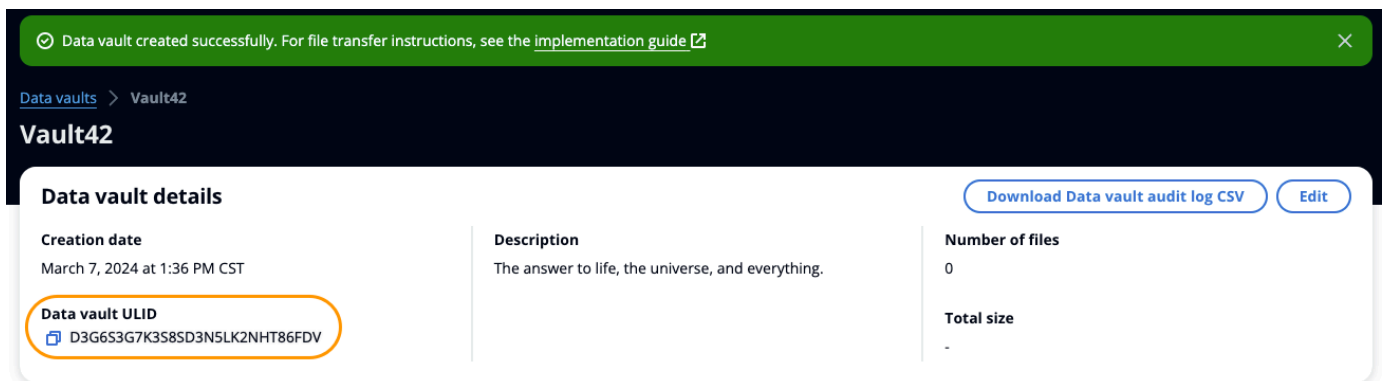
- [Create a data vault](#)
- [Create a source location](#)
- [Create a destination location and file transfer task](#)
- [Run the file transfer task in DEA](#)
- [Associate case files](#)
- [Disassociate case files](#)

Create a data vault

1. From the navigation, choose **Data vaults**.
2. Choose **Create data vault**.



3. For **Name**, enter a unique name for the data vault.
4. (Optional) Enter a description of the data vault.
5. Choose **Create**.
6. From the **Data vault details**, note the **Data vault ULID**. You will need this to create a DataSync destination location.



Create a source location

Using AWS DataSync, Digital Evidence Archive can support multiple storage types for data sources. For more information on supported source storage types, see [Where can I transfer my data with AWS DataSync?](#)

To set up your source location:

1. Follow the directions to [Create a source location for AWS DataSync](#).

2. Note your source location ARN if you intend to [the section called "Create a destination location and file transfer task using the DEA API"](#).

Note

If your source location is not an Amazon S3 bucket, the Lambda role that runs the file transfer tasks requires specific permissions to initiate the task. You must include the needed permissions in your stage file under the configuration variable `dataSyncSourcePermissions`.

Some common policies that you may need:

- `ec2:DescribeNetworkInterfaces`

Enables SMBs on an Amazon EC2 instance or Amazon EFS source locations.

- `fsx:DescribeFileSystems`

Enables Amazon FSx source locations.

Create a destination location and file transfer task

To transfer files into a data vault, you must create a destination location and file transfer task in AWS DataSync or using the DEA API. After the task is created, it can run in Digital Evidence Archive.

Create a destination location and file transfer task using AWS DataSync

These steps must be performed in the AWS DataSync console (<https://console.aws.amazon.com/datasync/>).

To create the destination location:

1. From the account that deployed DEA, find your `<stagename>-deamainstack` in AWS CloudFormation and note the available outputs.
2. From **Locations**, choose **Create location**.
3. Enter the following parameters for the destination location:

Parameter	Entry
Location type	Amazon S3
S3 bucket	<p>From the CloudFormation outputs, enter the data set bucket found using the key DeaBackendStackDeaS3Datasets74D75F63.</p> <p>If you don't see your data set bucket, then your AWS management account needs access to the bucket. You must set the ADMIN_ROLE_ARN and redeploy. For more information, see the section called "Step 3: Launch the solution".</p>
S3 storage class when used as a destination	Intelligent-Tiering
Folder	<p>/DATAVAULT{<i>dataVaultUlid</i>} /</p> <p>You may append additional paths to the folder path to specify a directory other than the root.</p> <p>For example: /DATAVAULT{dataVaultUlid}/Evidence/BodyCam/</p>
IAM role	From the CloudFormation outputs, enter the IAM role found using the key DeaBackendStackDeaDataSyncRoleF7B48276.

4. Choose **Create location**.

To create the file transfer task:

1. From **Tasks**, choose **Create task**.
2. For **Source location options**, select **Choose an existing location** and choose your source location from **Existing locations**.
3. Choose **Next**.

4. For **Destination location options**, select **Choose an existing location** and choose your destination location from **Existing locations**.
5. Choose **Next**.
6. On **Configure settings**, enter the following parameters:

Parameter	Entry
Task name	Enter a meaningful name for the task. This will appear in DEA.
Verification	Select Verify only the data transferred .
Overwrite files	Clear the check by Overwrite files .
Report type	Choose Standard report .
Report level	Choose Successes and errors .
S3 bucket for reports	From the CloudFormation outputs, enter the bucket ARN found using the key <code>DeaBackendStackDeaDataSyncReportsBucketNameB5D8462F</code> .
IAM role	From the CloudFormation outputs, enter the IAM role ARN found using the key <code>DeaBackendStackDeaDataSyncReportsRoleED5B265D</code> .
Logging	Choose the logging option which adheres to your organization's policies.

7. Choose **Next**.
8. Review your selections and choose **Create task**.
9. Note the task ULID before returning to the DEA user interface.

Create a destination location and file transfer task using the DEA API

Note

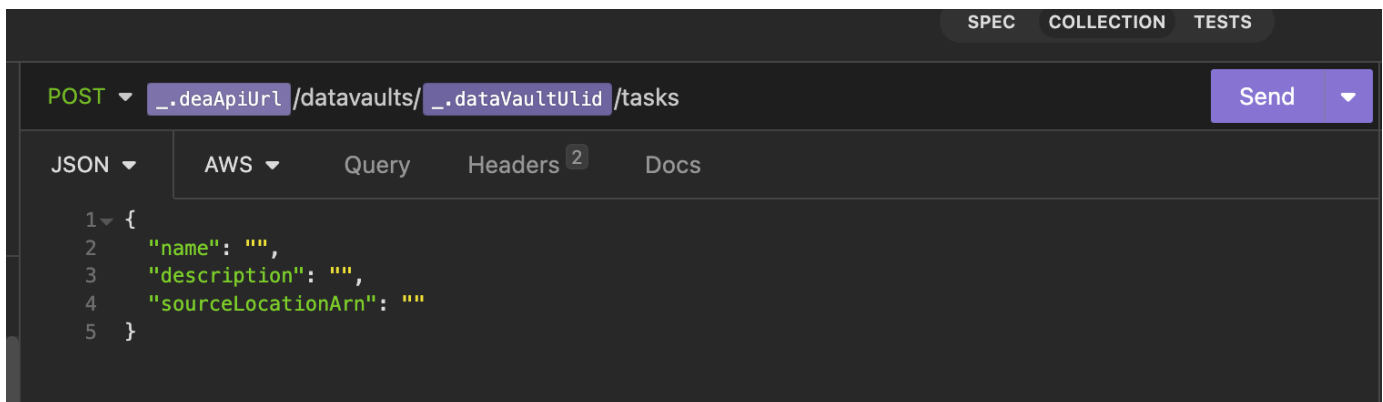
Before you begin, make sure that the DEA Lambda has access to create a task for the source location. Open your stage file and add your bucket ARN to the `dataSyncLocationBuckets` array. You must rebuild and redeploy following this change. Example: `"dataSyncLocationBuckets": ["arn:aws:s3:::bucket1", "arn:aws:s3:::bucket2"]`

1. *DOES THIS HAVE TO BE INSOMNIA? COULD THEY POTENTIALLY USE SOME OTHER CLIENT?*

From your API client, make a POST request to `/datavaults/{dataVaultUlid}/tasks`. The payload should include:

- name
- description
- sourceLocationArn

2. Enter the data vault ULID in your environment variables.

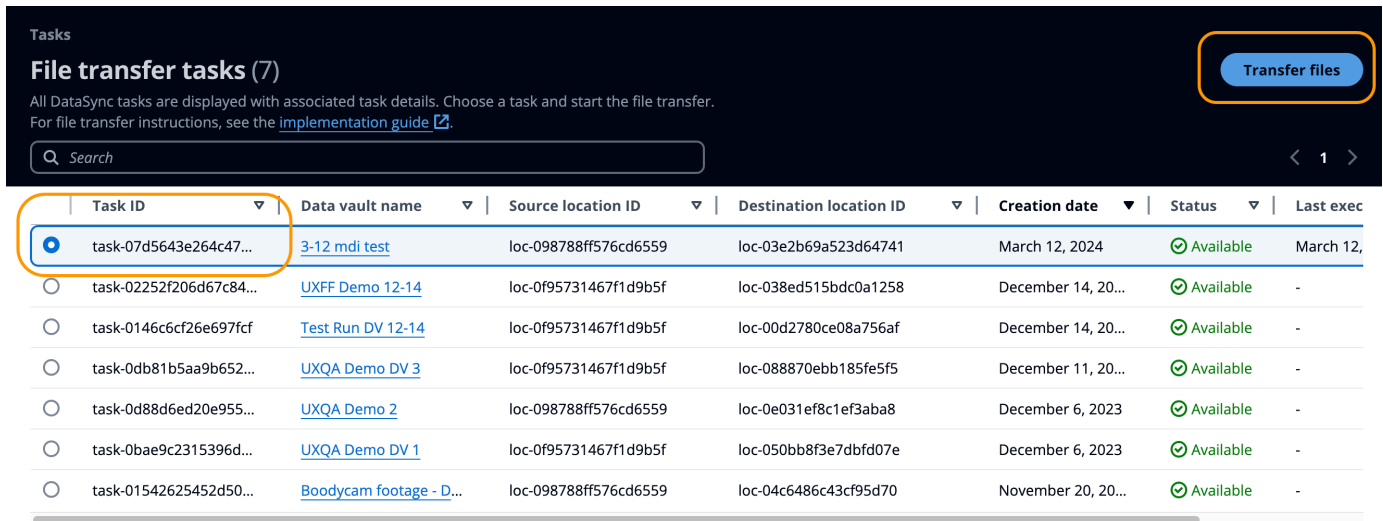


3. Note your task ID and return to the DEA user interface.

Run the file transfer task in DEA

From the DEA user interface, you can complete the mass data ingestion task you initiated in AWS DataSync.

1. From the navigation, choose **File transfer tasks**.
2. Use the task ID from DataSync to find the task you created.
3. Select your task and choose **Transfer files**.



Tasks

File transfer tasks (7)

All DataSync tasks are displayed with associated task details. Choose a task and start the file transfer. For file transfer instructions, see the [implementation guide](#).

Search

Task ID	Data vault name	Source location ID	Destination location ID	Creation date	Status	Last exec
task-07d5643e264c47...	3-12 mdi test	loc-098788ff576cd6559	loc-03e2b69a523d64741	March 12, 2024	Available	March 12,
task-02252f206d67c84...	UXFF Demo 12-14	loc-0f95731467f1d9b5f	loc-038ed515bdc0a1258	December 14, 20...	Available	-
task-0146c6cf26e697fcf	Test Run DV 12-14	loc-0f95731467f1d9b5f	loc-00d2780ce08a756af	December 14, 20...	Available	-
task-0db81b5aa9b652...	UXQA Demo DV 3	loc-0f95731467f1d9b5f	loc-088870ebb185fe5f5	December 11, 20...	Available	-
task-0d88d6ed20e955...	UXQA Demo 2	loc-098788ff576cd6559	loc-0e031ef8c1ef3aba8	December 6, 2023	Available	-
task-0bae9c2315396d...	UXQA Demo DV 1	loc-0f95731467f1d9b5f	loc-050bb8f3e7dbfd07e	December 6, 2023	Available	-
task-01542625452d50...	Boodycam footage - D...	loc-098788ff576cd6559	loc-04c6486c43cf95d70	November 20, 20...	Available	-

4. In the **Confirm details are correct** modal, review the transfer details and choose **Start file transfer**.

When the task completes, you will receive a notification confirming success. You can track the progress of your task in AWS DataSync. When the transfer completes successfully, you will see the transferred files in the data vault's detail page.

Associate case files

Once files are added to a data vault, they are not yet assigned to a case. You will need associate files to a case for case members to access the evidence.

Note

- For each association task, you can only associate up to 300 files per task.
- Files can only be downloaded from their associated cases, not from the data vault.

1. From the **Data vaults** page, open a data vault to view the details page.
2. Under **Files**, select the files or folders you want to associate with a case.
3. Choose **Associate to case**.

4. Choose the cases to which the files will be associated. You can associate files with more than one case.
5. Choose **Confirm**.

Disassociate case files

To disassociate files from a case, you must disassociate each file individually.

1. From the data vault where the file is stored, open the individual file's detail page.
2. Choose **Disassociate**.
3. Select the case(s) you want to disassociate the file from.
4. Choose **Disassociate**.

Developer guide

This developer guide provides additional instructions for a technical audience on how to customize and integrate with the Digital Evidence Archive on AWS solution. This includes information on accessing the Digital Evidence Archive source code and customer-facing API.

Source code

Visit our [GitHub repository](#) to download the source files for this solution and to share your customizations with others.

The Digital Evidence Archive templates are generated using the [AWS Cloud Development Kit \(AWS CDK\) \(AWS CDK\)](#). See the [README.md file](#) for additional information.

Integration guide

Adding Transport Layer Security (TLS) 1.2 support

For enhanced security, we recommend that you add Transport Layer Security (TLS) 1.2 support to your deployed Digital Evidence Archive on AWS solution. You can add TLS 1.2 support by setting a security policy in the Amazon API Gateway console, AWS Command Line Interface, or AWS SDK.

For more information on adding Transport Layer Security (TLS) 1.2 support, see [Choosing a minimum TLS version for a custom domain in API Gateway](#).

Customization guide

Change domain names post-deployment

You can use your organization's domain name with Digital Evidence Archive. If your organization has a personal domain name, you can add that domain to Digital Evidence Archive by assigning it in [Amazon Route 53](#).

Then, you can assign your organization's domain name to the current web address through [Amazon Route 53](#). Otherwise, you can register and host a new domain to use with Digital Evidence Archive through [Amazon Route 53](#).

Assign users to DEA

Note

For each user you want to have access to the DEA application, you must have the following fields defined for them:

- First name
- Last name
- Email
- DEARole

This custom attribute would be created during deployment.

Follow your IdP's documentation to give users their appropriate DEA role and assign them to the SAML 2.0 app integration you made during deployment. After this, these users should be able to log on to DEA using their existing credentials.

Log in to Digital Evidence Archive

1. Assign yourself to the DEA SAML 2.0 integration.
2. Make sure you have a DEARole attribute assigned to yourself, preferably a role that permits access to case APIs.
3. Get the UI URL from the launch outputs. (Its field name is DeaApiGatewayUiUrl).
4. Enter the URL into your browser, and you should be automatically redirected to your IdPs sign in page. If the button to **Create new case** is blue, then the IdP integration is complete.

Optional security settings

DEA Permissions Boundary

The CDK stack for DEA also creates a permissions boundary that blocks all access to protected DEA resources. AWS CloudTrail is enabled for DEA resources, so all access outside of DEA will be logged. However, for extra security you can attach the Permissions Boundary to all Console IAM Roles to block access to DEA resources.

1. Open the IAM console at <https://console.aws.amazon.com/iam/>.

You must use the admin account for your deployment.

2. Go to either:
 - **Users** to set boundaries around users.
 - **Roles** to set permissions around a role.
3. Choose a user or role.
4. Go to **Permissions Boundary**, and choose **set permissions boundary**.
5. In the search bar, enter **deaResourcesPermissionsBoundary** and choose it.
6. Choose **Set boundary**.

Repeat steps 2 through 6 for each User (or Role that users have access to).

Enabling AWS WAF

Digital Evidence Archive permits configuration of AWS WAF (AWS WAF) to protect the DEA-created Amazon API Gateway. For more information on how you can protect your deployment, see the following documentation:

- [Throttle requests for better throughput](#)
- [How do I apply a rate limit on a specific request parameter on URI in AWS WAF?](#)

Note

The default DEA limits are conservative, but you can increase the numbers alongside the Lambda service limit increases.

Antivirus

DEA does not protect against uploading viruses (as they can be considered evidence in certain investigations). Therefore, we recommend setting up antivirus and other security protections, such as CrowdStrike, to validate that your Criminal Justice devices stay secure.

API reference

The [API documentation for Digital Evidence Archive](#) provides comprehensive guidance on integrating with and leveraging DEA's features programmatically. It outlines endpoints for managing cases, files, users, and audit trails, offering detailed explanations of request and response formats. The documentation also includes code snippets and examples to assist developers in effectively using the API to store, retrieve, and manage digital evidence securely.

Reference

This section includes information about an optional feature for collecting unique metrics for this solution, pointers to related resources, and a list of builders who contributed to this solution.

Contributors

- Andrew Pecka
- Danielle Melvin
- Ajay Amuthan
- Daniel Wood
- Chase Chow
- Denis Estevez Quesada
- Miguel Abreu
- Kristen Caplan
- Brandi Hopkins

Notices

AWS' release, and your use, of this solution and any code provided thereunder (each a "solution") is subject to the AWS Customer Agreement or other applicable agreement between you and AWS governing your use of AWS Services, as well as the [AWS Shared Responsibility Model](#), which outlines your obligations for maintaining a secure environment. The solution is provided as an open-source building block that may assist you in accelerating the development of new or existing cloud capabilities. The open-source solution does not come with any standard SLA support. You are responsible for the end-to-end security and implementation of the solution code in your environment. Issues reported through GitHub will be responded to on a reasonable best-effort basis. AWS offers no guarantees regarding response times to submitted issues, requests for support, feature requests, or to the long-term roadmap for the solution. If you need deeper insight into the solution and/or the solution code, we suggest connecting with a qualified AWS Partner or AWS ProServe to gain additional clarity and resolve any challenges.

Customers are responsible for making their own independent assessment of the information in this document. This document: (a) is for informational purposes only, (b) represents AWS current

product offerings and practices, which are subject to change without notice, and (c) does not create any commitments or assurances from AWS and its affiliates, suppliers or licensors. AWS products or services are provided “as is” without warranties, representations, or conditions of any kind, whether express or implied. AWS responsibilities and liabilities to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

Digital Evidence Archive on AWS is licensed under the terms of the of the Apache License Version 2.0 available at [The Apache Software Foundation](https://www.apache.org/licenses/LICENSE-2.0).

Revisions

For more information, see the [CHANGELOG.md](#) file in the GitHub repository.

Date	Change
June 2023	Initial release
July 2023	Release version 1.0.1 - For more information, refer to the CHANGELOG.md file in the GitHub repository.
July 2023	Release version 1.0.2 - For more information, refer to the CHANGELOG.md file in the GitHub repository.
August 2023	Release version 1.0.3 - For more information, refer to the CHANGELOG.md file in the GitHub repository.
October 2023	Release version 1.0.4 - For more information, refer to the CHANGELOG.md file in the GitHub repository.
October 2023	Release version 1.0.5 - For more information, refer to the CHANGELOG.md file in the GitHub repository.
December 2023	Release version 1.0.6 - For more information, refer to the CHANGELOG.md file in the GitHub repository.
January 2024	Release version 1.0.7 - For more information, refer to the CHANGELOG.md file in the GitHub repository.

Date	Change
March 2024	Release version 1.1.0 - For more information, refer to the CHANGELOG.md file in the GitHub repository.