
Amazon CloudFront

Developer Guide



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

What is Amazon CloudFront?	1
How you set up CloudFront to deliver content	1
Use cases	3
Accelerate static website content delivery	3
Serve video on demand or live streaming video	4
Encrypt specific fields throughout system processing	4
Customize at the edge	4
Serve private content by using Lambda@Edge customizations	4
How CloudFront delivers content	5
How CloudFront delivers content to your users	5
How CloudFront works with regional edge caches	6
Locations and IP address ranges of CloudFront edge servers	7
Accessing CloudFront	7
How to get started with Amazon CloudFront	8
AWS Identity and Access Management	8
CloudFront pricing	8
Choosing the price class for a CloudFront distribution	10
Setting up	11
Sign up for AWS	11
Access your account	11
Access the console	12
Access the API, AWS CLI, AWS Tools for Windows PowerShell, or the AWS SDKs	12
Create an IAM user	12
Set up the AWS Command Line Interface or AWS Tools for Windows PowerShell	14
Download an AWS SDK	14
Getting started	15
Getting started with a simple distribution	15
Prerequisites	16
Step 1: Upload your content to Amazon S3 and grant object permissions	16
Step 2: Create a CloudFront distribution	17
Step 3: Test your links	18
Getting started with AWS for WordPress	19
Prerequisites	20
Step 1: Install the plugin	22
Step 2: Configure and use CloudFront with the plugin	22
(Optional) Deactivate site acceleration	24
(Optional) Remove site acceleration and delete the CloudFront distribution	25
(Optional) Deactivate and remove the plugin	25
(Optional) Create a CloudFront distribution for Amazon Polly content	26
Troubleshooting	26
Getting started with a secure static website	28
Solution overview	29
Deploying the solution	30
Working with distributions	33
Overview of distributions	33
Actions you can use with distributions	34
Required fields for creating and updating distributions	34
Creating, Updating, and Deleting Distributions	36
Steps for Creating a Distribution	36
Creating a Distribution	37
Values That You Specify	38
Values That Are Displayed	61
Testing a Distribution	62
Updating a Distribution	63

Tagging a Distribution	64
Deleting a Distribution	65
Using Different Origins	66
Using Amazon S3 Buckets for Your Origin	67
Using Amazon S3 Buckets Configured as Website Endpoints for Your Origin	68
Using a MediaStore Container or a MediaPackage Channel for Your Origin	68
Using Amazon EC2 or Other Custom Origins	68
Using Origin Groups	69
Adding CloudFront When You Have Amazon S3 Content	69
Moving an Amazon S3 Bucket to a Different Region	71
Using Custom URLs	71
Adding an Alternate Domain Name	72
Moving an Alternate Domain Name to a Different CloudFront Distribution	74
Removing an Alternate Domain Name	77
Using Wildcards in Alternate Domain Names That You Add to CloudFront	78
Requirements for Using Alternate Domain Names	78
Restrictions on Using Alternate Domain Names	80
Using Websockets	81
How the WebSocket Protocol Works	81
WebSocket Requirements	81
Working with policies	83
Controlling the cache key	83
Understanding cache policies	84
Creating cache policies	88
Using the managed cache policies	91
Understanding the cache key	93
Controlling origin requests	95
Understanding origin request policies	96
Creating origin request policies	97
Using the managed origin request policies	100
Using the CloudFront HTTP headers	102
Headers for determining the viewer's device type	102
Headers for determining the viewer's location	103
Other CloudFront headers	103
Adding, Removing, or Replacing Content	104
Adding and Accessing Content	104
Updating Existing Content	104
Updating Existing Files Using Versioned File Names	105
Updating Existing Content Using the Same File Names	105
Removing Content so CloudFront Won't Distribute It	106
Customizing File URLs	106
Using Your Own Domain Name (Example.com)	106
Using a Trailing Slash (/) in URLs	107
Creating Signed URLs for Restricted Content	107
Creating URLs for RTMP Distribution Media Files	107
Specifying a Default Root Object	107
How to Specify a Default Root Object	107
How Headers Work With Default Root Objects	108
How CloudFront Works if You Don't Define a Root Object	109
Invalidating Files	110
Choosing Between Invalidating Files and Using Versioned File Names	110
Determining Which Files to Invalidate	111
Specifying the Files to Invalidate	111
Invalidating Files Using the Console	113
Invalidating Files Using the CloudFront API	115
Concurrent Invalidation Request Maximum	116
Paying for File Invalidation	116

Serving compressed files	116
Configuring a CloudFront distribution to compress content	117
Using CloudFront to compress your content	117
Configuring Secure Access and Restricting Access to Content	121
Using HTTPS with CloudFront	121
Requiring HTTPS Between Viewers and CloudFront	122
Requiring HTTPS to a Custom Origin	123
Requiring HTTPS to an Amazon S3 Origin	127
Supported protocols and ciphers between viewers and CloudFront	128
Supported protocols and ciphers between CloudFront and the origin	130
Charges for HTTPS connections	132
Using Alternate Domain Names and HTTPS	132
Choosing How CloudFront Serves HTTPS Requests	132
Requirements for Using SSL/TLS Certificates with CloudFront	134
Quotas on Using SSL/TLS Certificates with CloudFront (HTTPS Between Viewers and CloudFront Only)	137
Configuring Alternate Domain Names and HTTPS	139
Determining the Size of the Public Key in an SSL/TLS Certificate	142
Increasing the Quotas for SSL/TLS Certificates	142
Rotating SSL/TLS Certificates	143
Reverting from a Custom SSL/TLS Certificate to the Default CloudFront Certificate	144
Switching from a Custom SSL/TLS Certificate with Dedicated IP Addresses to SNI	145
Restricting Content with Signed URLs and Signed Cookies	145
Overview of Serving Private Content	146
Task List: Serving Private Content	149
Specifying Your Trusted Signers	150
Choosing Between Signed URLs and Signed Cookies	156
Using Signed URLs	157
Using Signed Cookies	175
Using a Linux Command and OpenSSL for Base64-Encoding and Encryption	189
Code Examples for Signed URLs	189
Restricting Access to Amazon S3 Content	209
Overview of OAI Setup	210
Creating a CloudFront OAI and Adding it to Your Distribution	211
Granting the OAI Permission to Read Files in Your Amazon S3 Bucket	212
Using an OAI in Amazon S3 Regions that Support Only Signature Version 4 Authentication	215
Using AWS WAF to Control Access to Your Content	216
Geo-Restricting Content	217
Using CloudFront Geo Restriction	217
Using a Third-Party Geolocation Service	218
Using Field-Level Encryption to Help Protect Sensitive Data	219
Overview of Field-Level Encryption	221
Setting Up Field-Level Encryption	222
Decrypting Data Fields at Your Origin	225
Optimizing Content Caching and Availability	228
Caching with Edge Caches	228
Improving Your Cache Hit Ratio	228
Specifying How Long CloudFront Caches Your Objects	229
Caching Based on Query String Parameters	229
Caching Based on Cookie Values	229
Caching Based on Request Headers	230
Remove Accept-Encoding Header When Compression is Not Needed	231
Serving Media Content by Using HTTP	231
Increasing Availability with Origin Failover	231
Creating an Origin Group	232
Controlling Origin Timeouts and Attempts	233
Use Origin Failover with Lambda@Edge Functions	234

Use Custom Error Pages with Origin Failover	236
Managing Cache Expiration	236
Using Headers to Control Cache Duration for Individual Objects	237
Specifying the Amount of Time that CloudFront Caches Objects for Web Distributions	238
Specifying the Minimum Time that CloudFront Caches Objects for RTMP Distributions	240
Adding Headers to Your Objects Using the Amazon S3 Console	241
Caching and Query String Parameters	241
Console and API Settings for Query String Forwarding and Caching	243
Optimizing Caching	243
Query String Parameters and CloudFront Access Logs	244
Caching Content Based on Cookies	244
Caching Content Based on Request Headers	246
Headers and Distributions - Overview	247
Selecting the Headers to Base Caching On	247
Configuring CloudFront to Respect CORS Settings	248
Configuring Caching Based on the Device Type	248
Configuring Caching Based on the Language of the Viewer	249
Configuring Caching Based on the Location of the Viewer	249
Configuring Caching Based on the Protocol of the Request	249
Configuring Caching For Compressed Files	249
How Caching Based on Headers Affects Performance	249
How the Case of Headers and Header Values Affects Caching	250
Headers that CloudFront Returns to the Viewer	250
Troubleshooting	251
Troubleshooting Distribution Issues	251
CloudFront Returns an InvalidViewerCertificate Error When I Try to Add an Alternate Domain Name	251
I Can't View the Files in My Distribution	252
Error Message: Certificate: <certificate-id> Is Being Used by CloudFront	253
Troubleshooting Error Responses from Your Origin	254
HTTP 400 Status Code (Bad Request)	254
HTTP 500 Status Code (Lambda Execution Error)	255
HTTP 502 Status Code (Bad Gateway)	255
HTTP 502 Status Code (Lambda Validation Error)	258
HTTP 503 Status Code (Lambda Limit Exceeded)	258
HTTP 503 Status Code (Service Unavailable)	258
HTTP 504 Status Code (Gateway Timeout)	259
Load Testing CloudFront	262
Request and Response Behavior	263
Request and Response Behavior for Amazon S3 Origins	263
How CloudFront Processes HTTP and HTTPS Requests	263
How CloudFront Processes and Forwards Requests to Your Amazon S3 Origin Server	263
How CloudFront Processes Responses from Your Amazon S3 Origin Server	268
Request and Response Behavior for Custom Origins	269
How CloudFront Processes and Forwards Requests to Your Custom Origin Server	269
How CloudFront Processes Responses from Your Custom Origin Server	279
Request and Response Behavior for Origin Groups	282
Adding Custom Headers to Origin Requests	282
Use Cases for Origin Custom Headers	283
Configuring CloudFront to Add Custom Headers to Origin Requests	283
Custom Headers that CloudFront Can't Add to Origin Requests	284
Configure CloudFront to Forward Authorization Headers	284
How Range GETs Are Processed	285
How CloudFront Processes HTTP 3xx Status Codes from Your Origin	285
How CloudFront Processes and Caches HTTP 4xx and 5xx Status Codes from Your Origin	286
How CloudFront Processes Errors When You Have Configured Custom Error Pages	287
How CloudFront Processes Errors When You Have Not Configured Custom Error Pages	288

HTTP 4xx and 5xx Status Codes that CloudFront Caches	289
Generating Custom Error Responses	291
Creating a Custom Error Page for Specific HTTP Status Codes	291
Storing in Different Locations	293
Changing Response Codes	293
Controlling How Long Errors are Cached	294
How CloudFront Responds When a Custom Error Page Is Unavailable	294
Pricing for Custom Error Pages	295
Configuring Error Response Behavior	295
Video on Demand and Live Streaming Video	297
About Streaming Video: Video on Demand and Live Streaming	297
Delivering Video on Demand	298
Configuring Video on Demand for Microsoft Smooth Streaming	298
Delivering Live Streaming Video	300
Serving Video Using AWS Elemental MediaStore as the Origin	300
Serving Live Video Formatted with AWS Elemental MediaPackage	301
Working with RTMP Distributions	305
RTMP Distributions	305
How RTMP Distributions Work	306
Task List for Streaming Media Files Using RTMP	307
Creating an RTMP Distribution Using the CloudFront Console	308
Values that You Specify When You Create or Update an RTMP Distribution	309
Values that CloudFront Displays in the Console When You Create or Update an RTMP Distribution	313
Configuring the Media Player	314
Using an Amazon S3 Bucket as the Origin for an RTMP Distribution	315
Creating Multiple RTMP Distributions for an Origin Server	316
Restricting Access Using Crossdomain.xml	316
Error Codes for RTMP Distributions	316
Troubleshooting RTMP Distributions	316
Customizing with Lambda@Edge	318
Get Started Creating and Using Lambda@Edge Functions	319
Tutorial: Creating a Simple Function	320
Setting IAM Permissions and Roles	329
IAM Permissions Required to Associate Lambda Functions with CloudFront Distributions	330
Function Execution Role for Service Principals	330
Service-Linked Roles for Lambda@Edge	331
Writing and Creating Functions	335
Writing Functions for Lambda@Edge	335
Creating a Lambda@Edge Function in the Lambda Console	336
Editing a Lambda Function for Lambda@Edge	337
Creating Lambda Functions and CloudFront Triggers Programmatically	338
Adding Triggers	339
CloudFront Events That Can Trigger a Lambda Function	339
How to Decide Which CloudFront Event to Use to Trigger a Lambda Function	341
Adding Triggers by Using the Lambda Console	341
Adding Triggers by Using the CloudFront Console	342
Testing and Debugging	343
Testing your Lambda@Edge Functions	344
Identifying Lambda Function Errors in CloudFront	344
Troubleshooting Invalid Lambda Function Responses (Validation Errors)	349
Troubleshooting Lambda Function Execution Errors	350
Determining the Lambda@Edge Region	350
Determining if Your Account Pushes Logs to CloudWatch	351
CloudWatch Metrics and CloudWatch Logs	351
CloudWatch Metrics	351
CloudWatch Logs	351

Deleting Functions and Replicas	352
Event Structure	353
Dynamic Origin Selection	353
Request Events	353
Response Events	358
Working with Requests and Responses	363
Using Lambda@Edge Functions with Origin Failover	364
Generating HTTP Responses in Request Triggers	364
Updating HTTP Responses in Origin-Response Triggers	366
Accessing the Request Body by Choosing the Include Body Option	367
Example Functions	367
General Examples	368
Generating Responses - Examples	371
Working with Query Strings - Examples	373
Personalize Content by Country or Device Type Headers - Examples	377
Content-Based Dynamic Origin Selection - Examples	381
Updating Error Statuses - Examples	388
Accessing the Request Body - Examples	390
Requirements and Restrictions	393
CloudFront Distributions and Associations	394
CloudFront Triggers for Lambda Functions	394
CloudWatch Logs	394
Headers	394
HTTP Status Codes	396
Lambda Function Supported Runtimes and Configuration	397
Quotas	397
Microsoft Smooth Streaming	399
Network Access	399
Query String Parameters	399
Maximum Size for Body with the Include Body Option	399
Tagging	399
URI	400
URI and Query String Encoding	400
Reports, Metrics, and Logs	401
AWS Billing and Usage Reports for CloudFront	401
AWS Billing Report for CloudFront	401
AWS Usage Report for CloudFront	402
Interpreting Your AWS Bill and the AWS Usage Report for CloudFront	403
CloudFront Console Reports	405
CloudFront Cache Statistics Reports	407
CloudFront Popular Objects Report	411
CloudFront Top Referrers Report	414
CloudFront Usage Reports	416
CloudFront Viewers Reports	421
Monitoring CloudFront with Amazon CloudWatch	428
Viewing CloudFront and Lambda@Edge metrics	429
Setting alarms	432
Downloading data	433
Getting metrics using the API	435
CloudFront logging	438
Logging requests	438
Logging service activity	438
Using Standard Logs (Access Logs)	439
Real-time logs	457
Capturing API Requests with CloudTrail	466
Tracking Configuration Changes with AWS Config	472
Set up AWS Config with CloudFront	472

View CloudFront Configuration History	472
Security	474
Data Protection	474
Encryption in Transit	475
Encryption at Rest	475
Restrict Access to Content	475
Identity and Access Management	476
Authentication	476
Access Control	478
Overview of Managing Access	478
Using IAM Policies for CloudFront	483
CloudFront API Permissions Reference	487
Logging and Monitoring	493
Compliance Validation	493
CloudFront Compliance Best Practices	493
Resilience	495
CloudFront Origin Failover	495
Infrastructure Security	495
Quotas	496
General Quotas	496
General Quotas on Web Distributions	497
General Quotas on Policies	497
Quotas on WebSocket Connections	498
Quotas on Whitelisted Cookies (Web Distributions Only)	498
Quotas on Whitelisted Query Strings (Web Distributions Only)	499
Quotas on Custom Headers (Web Distributions Only)	499
Quotas on SSL Certificates (Web Distributions Only)	499
Quotas on Invalidations	500
Quotas on Field-Level Encryption	500
Quotas on Lambda@Edge	500
Size Quotas on URI and Query String	501
Size Quotas on Request Body with the Include Body Option	501
Request Timeout	502
Quotas on RTMP Distributions	502
Related Information	503
Additional Amazon CloudFront Documentation	503
Getting Support	503
CloudFront Developer Tools and SDKs	503
Tips from the Amazon Web Services Blog	504
Document history	505
AWS glossary	508

What is Amazon CloudFront?

Amazon CloudFront is a web service that speeds up distribution of your static and dynamic web content, such as .html, .css, .js, and image files, to your users. CloudFront delivers your content through a worldwide network of data centers called edge locations. When a user requests content that you're serving with CloudFront, the user is routed to the edge location that provides the lowest latency (time delay), so that content is delivered with the best possible performance.

- If the content is already in the edge location with the lowest latency, CloudFront delivers it immediately.
- If the content is not in that edge location, CloudFront retrieves it from an origin that you've defined—such as an Amazon S3 bucket, a MediaPackage channel, or an HTTP server (for example, a web server) that you have identified as the source for the definitive version of your content.

As an example, suppose that you're serving an image from a traditional web server, not from CloudFront. For example, you might serve an image, `sunsetphoto.png`, using the URL `http://example.com/sunsetphoto.png`.

Your users can easily navigate to this URL and see the image. But they probably don't know that their request was routed from one network to another—through the complex collection of interconnected networks that comprise the internet—until the image was found.

CloudFront speeds up the distribution of your content by routing each user request through the AWS backbone network to the edge location that can best serve your content. Typically, this is a CloudFront edge server that provides the fastest delivery to the viewer. Using the AWS network dramatically reduces the number of networks that your users' requests must pass through, which improves performance. Users get lower latency—the time it takes to load the first byte of the file—and higher data transfer rates.

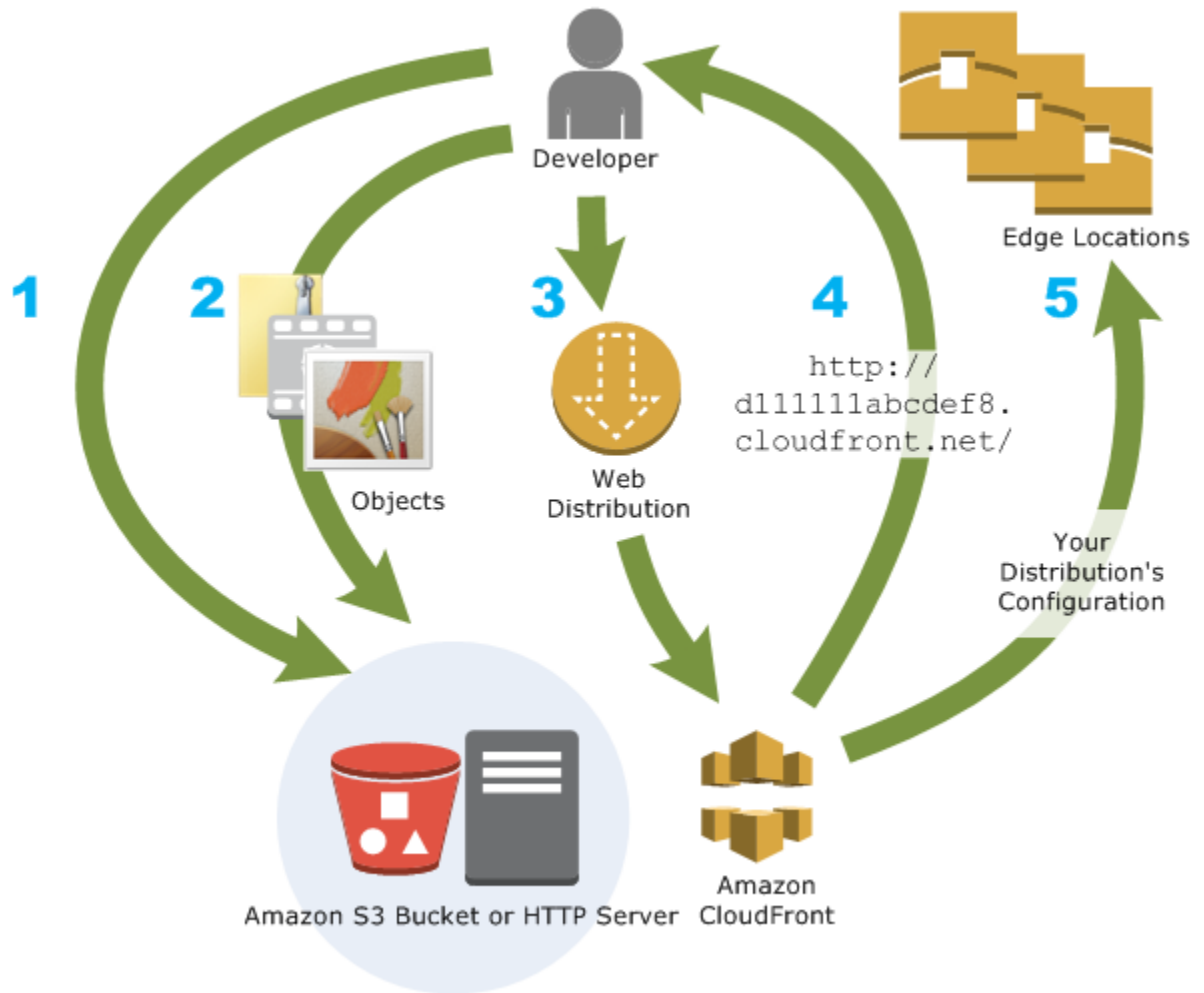
You also get increased reliability and availability because copies of your files (also known as *objects*) are now held (or cached) in multiple edge locations around the world.

Topics

- [How you set up CloudFront to deliver content \(p. 1\)](#)
- [CloudFront use cases \(p. 3\)](#)
- [How CloudFront delivers content \(p. 5\)](#)
- [Locations and IP address ranges of CloudFront edge servers \(p. 7\)](#)
- [Accessing CloudFront \(p. 7\)](#)
- [How to get started with Amazon CloudFront \(p. 8\)](#)
- [AWS Identity and Access Management \(p. 8\)](#)
- [CloudFront pricing \(p. 8\)](#)

How you set up CloudFront to deliver content

You create a CloudFront distribution to tell CloudFront where you want content to be delivered from, and the details about how to track and manage content delivery. Then CloudFront uses computers—edge servers—that are close to your viewers to deliver that content quickly when someone wants to see it or use it.



How you configure CloudFront to deliver your content

1. You specify *origin servers*, like an Amazon S3 bucket or your own HTTP server, from which CloudFront gets your files which will then be distributed from CloudFront edge locations all over the world.

An origin server stores the original, definitive version of your objects. If you're serving content over HTTP, your origin server is either an Amazon S3 bucket or an HTTP server, such as a web server. Your HTTP server can run on an Amazon Elastic Compute Cloud (Amazon EC2) instance or on a server that you manage; these servers are also known as *custom origins*.

If you use the Adobe Flash Media Server RTMP protocol to distribute media files on demand, your origin server is always an Amazon S3 bucket.

2. You upload your files to your origin servers. Your files, also known as *objects*, typically include web pages, images, and media files, but can be anything that can be served over HTTP or a supported version of Adobe RTMP, the protocol used by Adobe Flash Media Server.

If you're using an Amazon S3 bucket as an origin server, you can make the objects in your bucket publicly readable, so that anyone who knows the CloudFront URLs for your objects can access them. You also have the option of keeping objects private and controlling who accesses them. See [Serving Private Content with Signed URLs and Signed Cookies](#) (p. 145).

3. You create a CloudFront *distribution*, which tells CloudFront which origin servers to get your files from when users request the files through your web site or application. At the same time, you specify details such as whether you want CloudFront to log all requests and whether you want the distribution to be enabled as soon as it's created.
4. CloudFront assigns a domain name to your new distribution that you can see in the CloudFront console, or that is returned in the response to a programmatic request, for example, an API request. If you like, you can add an alternate domain name to use instead.
5. CloudFront sends your distribution's configuration (but not your content) to all of its *edge locations* or *points of presence* (POPs)—collections of servers in geographically-dispersed data centers where CloudFront caches copies of your files.

As you develop your website or application, you use the domain name that CloudFront provides for your URLs. For example, if CloudFront returns `d11111abcdef8.cloudfront.net` as the domain name for your distribution, the URL for `logo.jpg` in your Amazon S3 bucket (or in the root directory on an HTTP server) is `http://d11111abcdef8.cloudfront.net/logo.jpg`.

Or you can set up CloudFront to use your own domain name with your distribution. In that case, the URL might be `http://www.example.com/logo.jpg`.

Optionally, you can configure your origin server to add headers to the files, to indicate how long you want the files to stay in the cache in CloudFront edge locations. By default, each file stays in an edge location for 24 hours before it expires. The minimum expiration time is 0 seconds; there isn't a maximum expiration time. For more information, see [Managing How Long Content Stays in an Edge Cache \(Expiration\)](#) (p. 236).

CloudFront use cases

Using CloudFront can help you accomplish a variety of goals. This section lists just a few, together with links to more information, to give you an idea of the possibilities.

Topics

- [Accelerate static website content delivery](#) (p. 3)
- [Serve video on demand or live streaming video](#) (p. 4)
- [Encrypt specific fields throughout system processing](#) (p. 4)
- [Customize at the edge](#) (p. 4)
- [Serve private content by using Lambda@Edge customizations](#) (p. 4)

Accelerate static website content delivery

CloudFront can speed up the delivery of your static content (for example, images, style sheets, JavaScript, and so on) to viewers across the globe. By using CloudFront, you can take advantage of the AWS backbone network and CloudFront edge servers to give your viewers a fast, safe, and reliable experience when they visit your website.

A simple approach for storing and delivering static content is to use an Amazon S3 bucket. Using S3 together with CloudFront has a number of advantages, including the option to use Origin Access Identity (OAI) to easily restrict access to your S3 content.

For more information about using S3 together with CloudFront, including a AWS CloudFormation template to help you get started quickly, see [Amazon S3 + Amazon CloudFront: A Match Made in the Cloud](#).

Serve video on demand or live streaming video

CloudFront offers several options for streaming your media to global viewers—both pre-recorded files and live events.

- For video on demand (VOD) streaming, you can use CloudFront to stream in common formats such as MPEG DASH, Apple HLS, Microsoft Smooth Streaming, and CMAF, to any device.
- For broadcasting a live stream, you can cache media fragments at the edge, so that multiple requests for the manifest file that delivers the fragments in the right order can be combined, to reduce the load on your origin server.

For more information about how to deliver streaming content with CloudFront, see [Video on Demand and Live Streaming Video with CloudFront](#) (p. 297).

Encrypt specific fields throughout system processing

When you configure HTTPS with CloudFront, you already have secure end-to-end connections to origin servers. When you add field-level encryption, you can protect specific data throughout system processing in addition to HTTPS security, so that only certain applications at your origin can see the data.

To set up field-level encryption, you add a public key to CloudFront, and then specify the set of fields that you want to be encrypted with the key. For more information, see [Using Field-Level Encryption to Help Protect Sensitive Data](#) (p. 219).

Customize at the edge

Running serverless code at the edge opens up a number of possibilities for customizing the content and experience for viewers, at reduced latency. For example, you can return a custom error message when your origin server is down for maintenance, so viewers don't get a generic HTTP error message. Or you can use a function to help authorize users and control access to your content, before CloudFront forwards a request to your origin.

Using Lambda@Edge with CloudFront enables a variety of ways to customize the content that CloudFront delivers. To learn more about Lambda@Edge and how to create and deploy functions with CloudFront, see [Customizing Content at the Edge with Lambda@Edge](#) (p. 318). To see a number of code samples that you can customize for your own solutions, see [Lambda@Edge Example Functions](#) (p. 367).

Serve private content by using Lambda@Edge customizations

Using Lambda@Edge can help you configure your CloudFront distribution to serve private content from your own custom origin, as an option to using signed URLs or signed cookies.

You can use several techniques to restrict access to your origin exclusively to CloudFront, including using whitelisting CloudFront IPs in your firewall and using a custom header to carry a shared secret.

For more information and step-by-step instructions, including sample code, see [Serving Private Content Using Amazon CloudFront & AWS Lambda@Edge](#).

How CloudFront delivers content

After some initial setup, CloudFront works together with your website or application and speeds up delivery of your content. This section explains how CloudFront serves your content when viewers request it.

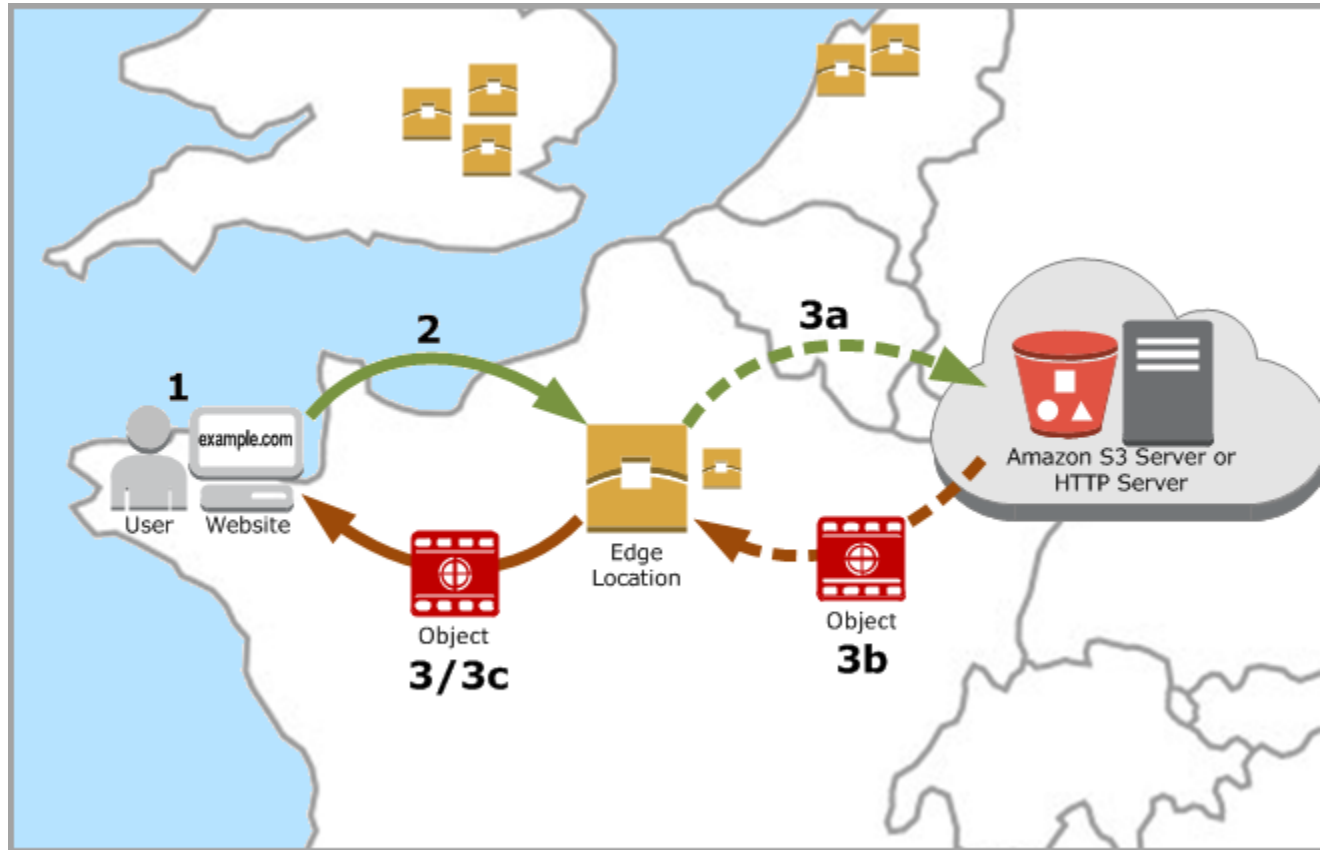
Topics

- [How CloudFront delivers content to your users \(p. 5\)](#)
- [How CloudFront works with regional edge caches \(p. 6\)](#)

How CloudFront delivers content to your users

After you configure CloudFront to deliver your content, here's what happens when users request your files:

1. A user accesses your website or application and requests one or more files, such as an image file and an HTML file.
2. DNS routes the request to the CloudFront POP (edge location) that can best serve the request—typically the nearest CloudFront POP in terms of latency—and routes the request to that edge location.
3. In the POP, CloudFront checks its cache for the requested files. If the files are in the cache, CloudFront returns them to the user. If the files are *not* in the cache, it does the following:
 - a. CloudFront compares the request with the specifications in your distribution and forwards the request for the files to your origin server for the corresponding file type—for example, to your Amazon S3 bucket for image files and to your HTTP server for HTML files.
 - b. The origin servers send the files back to the edge location.
 - c. As soon as the first byte arrives from the origin, CloudFront begins to forward the files to the user. CloudFront also adds the files to the cache in the edge location for the next time someone requests those files.



How CloudFront works with regional edge caches

CloudFront points of presence (POPs) (edge locations) make sure that popular content can be served quickly to your viewers. CloudFront also has *regional edge caches* that bring more of your content closer to your viewers, even when the content is not popular enough to stay at a POP, to help improve performance for that content.

Regional edge caches help with all types of content, particularly content that tends to become less popular over time. Examples include user-generated content, such as video, photos, or artwork; e-commerce assets such as product photos and videos; and news and event-related content that might suddenly find new popularity.

How regional caches work

Regional edge caches are CloudFront locations that are deployed globally, close to your viewers. They're located between your origin server and the POPs—global edge locations that serve content directly to viewers. As objects become less popular, individual POPs might remove those objects to make room for more popular content. Regional edge caches have a larger cache than an individual POP, so objects remain in the cache longer at the nearest regional edge cache location. This helps keep more of your content closer to your viewers, reducing the need for CloudFront to go back to your origin server, and improving overall performance for viewers.

When a viewer makes a request on your website or through your application, DNS routes the request to the POP that can best serve the user's request. This location is typically the nearest CloudFront edge location in terms of latency. In the POP, CloudFront checks its cache for the requested files. If the files are in the cache, CloudFront returns them to the user. If the files are not in the cache, the POPs go to the nearest regional edge cache to fetch the object.

In the regional edge cache location, CloudFront again checks its cache for the requested files. If the files are in the cache, CloudFront forwards the files to the POP that requested them. As soon as the first byte arrives from regional edge cache location, CloudFront begins to forward the files to the user. CloudFront also adds the files to the cache in the POP for the next time someone requests those files.

For files not cached at either the POP or the regional edge cache location, CloudFront compares the request with the specifications in your distributions and forwards the request for your files to the origin server. After your origin server sends the files back to the regional edge cache location, they are forwarded to the POP, and CloudFront forwards the files to the user. In this case, CloudFront also adds the files to the cache in the regional edge cache location in addition to the POP for the next time a viewer requests those files. This makes sure that all of the POPs in a region share a local cache, eliminating multiple requests to origin servers. CloudFront also keeps persistent connections with origin servers so files are fetched from the origins as quickly as possible.

Note

- Regional edge caches have feature parity with POPs. For example, a cache invalidation request removes an object from both POP caches and regional edge caches before it expires. The next time a viewer requests the object, CloudFront returns to the origin to fetch the latest version of the object.
- Proxy methods `PUT/POST/PATCH/OPTIONS/DELETE` go directly to the origin from the POPs and do not proxy through the regional edge caches.
- Dynamic requests, as determined at request time, do not flow through regional edge caches, but go directly to the origin.

Locations and IP address ranges of CloudFront edge servers

For a list of the locations of CloudFront edge servers, see the [Amazon CloudFront Product Details page](#).

Amazon Web Services (AWS) publishes its current IP address ranges in JSON format. To view the current ranges, download [ip-ranges.json](#). For more information, see [AWS IP Address Ranges](#) in the *Amazon Web Services General Reference*.

To find the IP address ranges that are associated with CloudFront edge servers, search `ip-ranges.json` for the following string:

```
"service": "CLOUDFRONT"
```

Alternatively, you can view only the CloudFront IP ranges [here](#).

Accessing CloudFront

You can access Amazon CloudFront in the following ways:

- **AWS Management Console** – The procedures throughout this guide explain how to use the AWS Management Console to perform tasks.
- **AWS SDKs** – If you're using a programming language that AWS provides an SDK for, you can use an SDK to access CloudFront. SDKs simplify authentication, integrate easily with your development environment, and provide access to CloudFront commands. For more information, see [Tools for Amazon Web Services](#).

- **CloudFront API** – If you're using a programming language that an SDK isn't available for, see the [Amazon CloudFront API Reference](#) for information about API actions and about how to make API requests.
- **AWS Command Line Interface** – For more information, see [Getting Set Up with the AWS Command Line Interface](#) in the *AWS Command Line Interface User Guide*.
- **AWS Tools for Windows PowerShell** – For more information, see [Setting up the AWS Tools for Windows PowerShell](#) in the *AWS Tools for Windows PowerShell User Guide*.

How to get started with Amazon CloudFront

For information about getting started with Amazon CloudFront, see the following topics in this guide:

- [Setting up Amazon CloudFront \(p. 11\)](#), which explains how to sign up for AWS, how to secure access to your AWS account, and how to set up programmatic access to CloudFront.
- [Getting started with Amazon CloudFront \(p. 15\)](#), which describes how to create a distribution that can serve content to viewers from your origin, like an Amazon S3 bucket or a website, and then test that it works.

AWS Identity and Access Management

Amazon CloudFront integrates with AWS Identity and Access Management (IAM), a service that lets your organization do the following:

- Create users and groups under your organization's AWS account
- Easily share your AWS account resources among the users in the account
- Assign unique security credentials to each user
- Granularly control user access to services and resources

For example, you can use IAM with CloudFront to control which users in your AWS account can create a new distribution or update cache behavior settings.

For general information about IAM, see the following:


- [Identity and Access Management in CloudFront \(p. 476\)](#)
- [Identity and Access Management \(IAM\)](#)
- [IAM User Guide](#)

CloudFront pricing

Amazon CloudFront is designed so you don't have to pay any up-front fees or commit to how much content you'll have. As with the other AWS services, you pay as you go and pay only for what you use. For more information, see [Amazon CloudFront Pricing](#).

AWS provides two usage reports for CloudFront: a billing report and a report that summarizes usage activity. To learn more about these reports, see [AWS Billing and Usage Reports for CloudFront \(p. 401\)](#).

The following diagram and list summarize the charges to use CloudFront.



Your monthly bill from AWS allocates your usage and dollar amounts by AWS service and function. The following explains the charges that are illustrated in the previous graphic. For more information, see [Amazon CloudFront Pricing](#).

1. **Charge for storage in an Amazon S3 bucket.** You pay normal Amazon S3 storage charges to store objects in your bucket. The charges appear in the Amazon S3 portion of your AWS statement.
2. **Charge for serving objects from edge locations.** You incur CloudFront charges when CloudFront responds to requests for your objects. The charges include data transfer for WebSocket data from server to client. The CloudFront charges appear in the CloudFront portion of your AWS statement as *region* -DataTransfer-Out-Bytes.
3. **Charge for submitting data to your origin.** You incur CloudFront charges when users transfer data to your origin, which includes DELETE, OPTIONS, PATCH, POST, and PUT requests. The charges include data transfer for WebSocket data from client to server. The CloudFront charges appear in the CloudFront portion of your AWS statement as *region* -DataTransfer-Out-OBytes.

Be aware of the following:

- You also incur a surcharge for HTTPS requests, and an additional surcharge for requests that also have field-level encryption enabled. For more information, see [Amazon CloudFront Pricing](#).

- You do not incur any additional CloudFront charges when you use origin groups. You continue to pay the same request fees and data transfer rates as you do when you use CloudFront with any other AWS or non-AWS origin. For more information, see [Using CloudFront Origin Groups \(p. 69\)](#).

Choosing the price class for a CloudFront distribution

CloudFront has edge locations all over the world. Our cost for each edge location varies and, as a result, the price that we charge you varies depending on the edge location from which CloudFront serves your requests.

CloudFront edge locations are grouped into geographic regions, and we've grouped regions into price classes. The default price class includes all regions. Another price class includes most regions (the United States; Canada; Europe; Hong Kong, Philippines, South Korea, Taiwan, and Singapore; Japan; India; South Africa; and Middle East regions) but excludes the most expensive regions. A third price class includes only the least expensive regions (the United States, Canada, and Europe regions).

By default, CloudFront responds to requests for your objects based only on performance: objects are served from the edge location for which latency is lowest for that viewer. If you're willing to accept higher latency for your viewers in some geographic regions in return for lower cost, you can choose a price class that doesn't include all CloudFront regions. Although CloudFront will serve your objects only from the edge locations in that price class, it still serves content from the edge location that has the lowest latency among the edge locations in your selected price class. However, some of your viewers, especially those in geographic regions that are not in your price class, may see higher latency than if your content were being served from all CloudFront edge locations. For example, if you choose the price class that includes only the United States and Europe, viewers in Australia and in Asia may experience higher latency than if you choose the price class that includes Australia and Asia.

If you choose a price class that does not include all edge locations, CloudFront may still occasionally serve requests for your content from an edge location in a region that is not included in your price class. When this happens, you are not charged the rate for the more expensive region from which your objects were served. Instead, you're charged the rate for the least expensive region in your selected price class.

You can choose a price class when you create or update a CloudFront distribution. For more information, see [Working with distributions \(p. 33\)](#).

If you're creating or updating a distribution by using the CloudFront API, one of the AWS SDKs, or AWS CloudFormation, see [DistributionConfig Complex Type](#) (search for `PriceClass`).

For more information about CloudFront pricing and price classes, see [Amazon CloudFront Pricing](#).

Setting up Amazon CloudFront

The overview and procedures in this section help you get started with AWS.

Topics

- [Sign up for AWS](#) (p. 11)
- [Access your account](#) (p. 11)
- [Create an IAM user](#) (p. 12)
- [Set up the AWS Command Line Interface or AWS Tools for Windows PowerShell](#) (p. 14)
- [Download an AWS SDK](#) (p. 14)

Sign up for AWS

When you sign up for AWS, your AWS account is automatically signed up for all services in AWS, including Amazon CloudFront. You are charged only for the services that you use.

If you have an AWS account already, skip to [Access your account](#) (p. 11). Otherwise, create one.

To create 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.

Note your AWS account number, because you'll need it later.

Tip

If you plan to use CloudFront to distribute content that you store in an S3 bucket, make sure that you also complete the steps to sign up for S3. For more information, see [Sign Up for Amazon S3](#).

Access your account

You use AWS services by using any of the following options:

- AWS Management Console
- API for each service
- AWS Command Line Interface (AWS CLI)
- AWS Tools for Windows PowerShell
- AWS SDKs

For each of those options, you need to access your AWS account by providing credentials that verify that you have permissions to use the services.

Access the console

To access the AWS Management Console for the first time, you provide an email address and a password. This combination of your email address and password is called your *root identity* or *root account credentials*. After you access your account for the first time, we strongly recommend that you don't use your root account credentials again for everyday use. Instead, you should create new credentials by using [AWS Identity and Access Management](#). To do that, you create a user account for yourself known as an *IAM user*, and then add the IAM user to an IAM group with administrative permissions or grant the IAM user administrative permissions. You then can access AWS using a special URL and the credentials for the IAM user. You also can add other IAM users later, and restrict their access to specified resources in the account.

Note

Some ad-blocking plugins for web browsers interfere with Amazon CloudFront console operations, which can cause the console to behave unpredictably. If you installed an ad-blocking plugin for your browser, we recommend that you add the URL for the CloudFront console, <https://console.aws.amazon.com/cloudfront/home>, to the whitelist for the plugin.

Access the API, AWS CLI, AWS Tools for Windows PowerShell, or the AWS SDKs

To use the API, the AWS CLI, AWS Tools for Windows PowerShell, or the AWS SDKs, you must create *access keys*. These keys consist of an access key ID and secret access key, which are used to sign programmatic requests that you make to AWS.

To create the keys, you sign in to the AWS Management Console. We strongly recommend that you sign in with your IAM user credentials instead of your root credentials. For more information, see [Managing Access Keys for IAM Users](#) in the *IAM User Guide*.

Create an IAM user

Use the following procedures to create a group for administrators, create an IAM user, and then add the IAM user to the administrators group. If you signed up for AWS but have not created an IAM user for yourself, you can create one using the IAM console. If you aren't familiar with using the console, see [Working with the AWS Management Console](#) for an overview.

To create an administrator user for yourself and add the user to an administrators group (console)

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

Note

We strongly recommend that you adhere to the best practice of using the **Administrator** IAM user below and securely lock away the root user credentials. Sign in as the root user only to perform a few [account and service management tasks](#).

2. In the navigation pane, choose **Users** and then choose **Add user**.
3. For **User name**, enter **Administrator**.
4. Select the check box next to **AWS Management Console access**. Then select **Custom password**, and then enter your new password in the text box.
5. (Optional) By default, AWS requires the new user to create a new password when first signing in. You can clear the check box next to **User must create a new password at next sign-in** to allow the new user to reset their password after they sign in.

6. Choose **Next: Permissions**.
7. Under **Set permissions**, choose **Add user to group**.
8. Choose **Create group**.
9. In the **Create group** dialog box, for **Group name** enter **Administrators**.
10. Choose **Filter policies**, and then select **AWS managed -job function** to filter the table contents.
11. In the policy list, select the check box for **AdministratorAccess**. Then choose **Create group**.

Note

You must activate IAM user and role access to Billing before you can use the `AdministratorAccess` permissions to access the AWS Billing and Cost Management console. To do this, follow the instructions in [step 1 of the tutorial about delegating access to the billing console](#).

12. Back in the list of groups, select the check box for your new group. Choose **Refresh** if necessary to see the group in the list.
13. Choose **Next: Tags**.
14. (Optional) Add metadata to the user by attaching tags as key-value pairs. For more information about using tags in IAM, see [Tagging IAM Entities](#) in the *IAM User Guide*.
15. Choose **Next: Review** to see the list of group memberships to be added to the new user. When you are ready to proceed, choose **Create user**.

You can use this same process to create more groups and users and to give your users access to your AWS account resources. To learn about using policies that restrict user permissions to specific AWS resources, see [Access Management](#) and [Example Policies](#).

To sign in as your new IAM user

1. Sign out of the AWS console.
2. Sign in by using the following URL, where `your_aws_account_id` is your AWS account number without the hyphens. For example, if your AWS account number is 1234-5678-9012, your AWS account ID is 123456789012:

```
https://your_aws_account_id.signin.aws.amazon.com/console/
```

3. Enter the IAM user name (not your email address) and password that you just created. When you're signed in, the navigation bar displays "`your_user_name @ your_aws_account_id`".

If you don't want the URL for your sign-in page to contain your AWS account ID, you can create an account alias.

To create an account alias and conceal your account ID

1. On the IAM console, choose **Dashboard** in the navigation pane.
2. On the dashboard, choose **Customize** and enter an alias such as your company name.
3. Sign out of the AWS console.
4. Sign in by using the following URL:

```
https://your_account_alias.signin.aws.amazon.com/console/
```

To verify the sign-in link for IAM users for your account, open the IAM console and check under **IAM users sign-in link** on the dashboard.

For more information about using IAM, see [Identity and Access Management in CloudFront \(p. 476\)](#).

Set up the AWS Command Line Interface or AWS Tools for Windows PowerShell

The AWS Command Line Interface (AWS CLI) is a unified tool for managing AWS services. For information about how to install and configure the AWS CLI, see [Getting Set Up with the AWS Command Line Interface](#) in the *AWS Command Line Interface User Guide*.

If you have experience with Windows PowerShell, you might prefer to use AWS Tools for Windows PowerShell. For more information, see [Setting up the AWS Tools for Windows PowerShell](#) in the *AWS Tools for Windows PowerShell User Guide*.

Download an AWS SDK

If you're using a programming language that AWS provides an SDK for, we recommend that you use an SDK instead of the Amazon CloudFront API. The SDKs make authentication simpler, integrate easily with your development environment, and provide easy access to CloudFront commands. For more information, see [Tools to Build on AWS](#).

Getting started with Amazon CloudFront

Get started with the basic steps to deliver your content with CloudFront by creating a simple CloudFront distribution or by using the AWS for WordPress plugin. If you already have a WordPress website, we recommend using the AWS for WordPress plugin to create a CloudFront distribution. If you aren't using WordPress, get started with a simple CloudFront distribution.

With both of the following options, you can get started for free using the [AWS Free Tier](#). When your 12-month free usage term expires or your use exceeds the tiers, you pay standard pay-as-you-go rates for the services that you use.

Topics

- [Getting started with a simple CloudFront distribution \(p. 15\)](#)
- [Getting started with the AWS for WordPress plugin \(p. 19\)](#)
- [Getting started with a secure static website \(p. 28\)](#)

Getting started with a simple CloudFront distribution

The procedures in this section show you how to use CloudFront to set up a basic configuration that does the following:

- Stores the original versions of your objects in one Amazon Simple Storage Service (Amazon S3) bucket
- Distributes content such as text or graphics
- Makes your objects accessible to everyone
- Uses the CloudFront domain name in URLs for your objects (for example, `http://d111111abcdef8.cloudfront.net/image.jpg`)
- Keeps your objects in CloudFront edge locations for the default duration of 24 hours (the minimum duration is 0 seconds)

Most of these options are customizable. For example, you can store your content on your own web server instead of using an S3 bucket, and you can restrict who has access to the content by using signed URLs or cookies. For information about how to customize your CloudFront distribution options, see [Steps for Creating a Distribution \(Overview\) \(p. 36\)](#).

You have to complete only a few basic steps to start delivering your content with CloudFront. The first step is signing up. After that, you create a CloudFront distribution, and then use the CloudFront domain name in URLs in your webpages or applications to reference the content.

Topics

- [Prerequisites \(p. 16\)](#)
- [Step 1: Upload your content to Amazon S3 and grant object permissions \(p. 16\)](#)
- [Step 2: Create a CloudFront distribution \(p. 17\)](#)
- [Step 3: Test your links \(p. 18\)](#)

Prerequisites

Before you begin, make sure that you've completed the steps in [Setting up Amazon CloudFront](#) (p. 11).

Step 1: Upload your content to Amazon S3 and grant object permissions

An Amazon S3 bucket is a container that can contain files (objects) or folders. CloudFront can distribute almost any type of file for you using an Amazon S3 bucket as the source. For example, CloudFront can distribute text, images, and videos. You can create multiple buckets, and there is no maximum for the amount of data that you can store on Amazon S3.

By default, your Amazon S3 bucket and all the files in it are private—only the AWS account that created the bucket has read/write permission to the files. If you want to allow anyone to access the files in your Amazon S3 bucket using CloudFront URLs, you must grant public read permissions to the objects. (This is one of the most common mistakes when working with CloudFront and Amazon S3. You must explicitly grant permissions to each object in an Amazon S3 bucket.)

Note

If you want to restrict who can download your content, you can use the CloudFront private content feature. For more information about distributing private content, see [Serving Private Content with Signed URLs and Signed Cookies](#) (p. 145).

To upload your content to Amazon S3 and grant read permissions to everyone

1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. On the Amazon S3 console, choose **Create bucket**.
3. In the **Create bucket** dialog box, on the **Name and region** page, do the following:
 - a. Enter a bucket name.

Important

For your bucket to work with CloudFront, the name must conform to DNS naming requirements. For more information, see [Bucket Restrictions and Limitations](#) in the *Amazon Simple Storage Service Developer Guide*.
 - b. For **Region**, choose an AWS Region for your bucket. By default, Amazon S3 creates buckets in the US East (N. Virginia) Region. We recommend that you choose a Region close to you to optimize latency and minimize costs, or you might choose another Region to address regulatory requirements.
4. Choose **Next**.
5. On the **Configure options** page, choose options for versioning, tagging, and other features.
6. Choose **Next**.
7. On the **Set permissions** page, clear the following check box:
 - **Block all public access**

You must allow public read access to the bucket and files so that CloudFront URLs can serve content from the bucket. However, you can restrict access to specific content by using the CloudFront private content feature. For more information, see [Serving Private Content with Signed URLs and Signed Cookies](#) (p. 145).

Select the check box for **I acknowledge that the current settings may result in this bucket and the objects within becoming public**.

8. Choose **Next**, and then choose **Create bucket**.

9. In the **Buckets** pane, choose your bucket, and then choose **Upload**.
10. On the **Select files** page, drag and drop your files to the bucket. Or choose **Add files**, and then choose the files that you want to upload.
11. Choose **Next**.
12. On the **Set permissions** page, for **Manage public permissions**, choose **Grant public read access to this object(s)**.
13. Choose **Next**.
14. Set any properties that you want for the object, such as encryption or tagging, and then choose **Next**.
15. Choose **Upload**.

After the upload is complete, you can navigate to the item by using its URL. For example:

```
https://<bucket name>.s3-<AWS Region>.amazonaws.com/<object name>
```

Use your Amazon S3 URL to verify that your content is publicly accessible, but remember that this is not the URL you'll use when you're ready to distribute your content with CloudFront.

Step 2: Create a CloudFront distribution

To create a CloudFront distribution

1. Open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
2. Choose **Create Distribution**.
3. On the **Select a delivery method for your content** page, in the **Web** section, choose **Get Started**.
4. On the **Create Distribution** page, for **Origin Domain Name**, choose the Amazon S3 bucket that you created earlier.

For the other settings under **Origin Settings**, accept the default values.

5. For the settings under **Default Cache Behavior Settings**, accept the default values.

CloudFront will:

- Forward all requests that use the CloudFront URL for your distribution (for example, `https://d1111111abcdef8.cloudfront.net/image.jpg`) to the Amazon S3 bucket that you specified in step 4 of this procedure.
- Allow end users to use either HTTP or HTTPS to access your objects.
- Respond to requests for your objects.
- Attempt to cache your objects at CloudFront edge locations for 24 hours, by default.
- Forward only the default request headers to your origin and not cache your objects based on the values in the headers.
- Exclude cookies and query string parameters, if any, when forwarding requests for objects to your origin. (Amazon S3 doesn't process cookies, and processes only some query string parameters.)
- Allow everyone to view your content.

CloudFront will not:

- Be configured to distribute media files in the Microsoft Smooth Streaming format.
- Automatically compress your content.

For more information about cache behavior options, see [Cache Behavior Settings \(p. 46\)](#).

- Under **Distribution Settings**, choose the values for your distribution:

Price Class

Choose the price class that corresponds with the maximum price that you want to pay for the CloudFront service. By default, CloudFront serves your objects from edge locations in all CloudFront Regions.

For more information about price classes and how your choice of price class affects CloudFront performance for your distribution, see [Choosing the price class for a CloudFront distribution \(p. 10\)](#). For information about CloudFront pricing, including how price classes map to CloudFront Regions, see [Amazon CloudFront Pricing](#).

AWS WAF Web ACL

Accept the default value, **None**.

Alternate Domain Names (CNAMEs)

Accept the default value—that is, leave the field empty.

SSL Certificate

Accept the default value, **Default CloudFront Certificate (*.cloudfront.net)**.

Supported HTTP Versions

Accept the default value, **HTTP/2, HTTP/1.1, HTTP/1.0**.

Default Root Object (Optional)

The object that you want CloudFront to request from your origin (for example, `index.html`) when a viewer requests the root URL of your distribution (`https://d111111abcdef8.cloudfront.net`) instead of an object in your distribution (`https://d111111abcdef8.cloudfront.net/product-description.html`). Specifying a default root object avoids exposing the contents of your distribution.

Logging (Optional)

Accept the default value, **Off**.

Enable IPv6

Accept the default value.

Comment (Optional)

Enter any comments that you want to save with the distribution.

Distribution State

Choose **Enabled**.

- Choose **Create Distribution**.
- After CloudFront creates your distribution, the value of the **Status** column for your distribution changes from **In Progress** to **Deployed**. This typically takes a few minutes.

The domain name that CloudFront assigns to your distribution appears in the list of distributions. (It also appears on the **General** tab for a selected distribution.)

Step 3: Test your links

After you've created your distribution, CloudFront knows where your Amazon S3 origin server is, and you know the domain name associated with the distribution. You can create a link to your Amazon S3 bucket content with that domain name and have CloudFront serve it.

Note

You must wait until the status of your distribution changes to **Deployed** before testing your links.

To link to your objects

1. Copy the following HTML into a new file:
 - Replace `<domain name>` with the domain name that CloudFront assigned to your distribution.
 - Replace `<object name>` with the name of a file in your Amazon S3 bucket.

```
<html>
<head>My CloudFront Test</head>
<body>
<p>My text content goes here.</p>
<p></p>
</body>
</html>
```

For example, if your domain name was `d111111abcdef8.cloudfront.net` and your object was `image.jpg`, the URL for the link would be this:

```
https://d111111abcdef8.cloudfront.net/image.jpg.
```

If your object is in a folder within your bucket, include the folder in the URL. For example, if `image.jpg` is located in an `images` folder, the URL would be this:

```
https://d111111abcdef8.cloudfront.net/images/image.jpg
```

2. Save the text in a file that has a `.html` file name extension.
3. Open your webpage in a browser to ensure that you can see your content. If you can't see the content, confirm that you performed all the steps correctly. You can also see the tips in [Troubleshooting \(p. 251\)](#).

The browser returns your page with the embedded image file, served from the edge location that CloudFront determined was appropriate to serve the object.

For more information about using CloudFront, see [Amazon CloudFront Related Information \(p. 503\)](#).

Getting started with the AWS for WordPress plugin

With the AWS for WordPress plugin, you can set up several AWS services, including Amazon CloudFront. With CloudFront, you can provide visitors to your WordPress website an accelerated viewing experience with content cached in edge locations around the world. When visitors come to your website, CloudFront routes them to the edge location that provides the lowest latency for a faster, more reliable experience.

The AWS for WordPress plugin creates a CloudFront distribution that is optimized for WordPress websites, using multiple cache behaviors to handle the different types of content on your website. The CloudFront features of the plugin work with websites hosted on [WordPress.com](#) and with self-hosted WordPress websites on Amazon Lightsail, Amazon EC2, or another web hosting platform.

You can also use the plugin to set up other AWS services like [Amazon Polly](#) and [Amazon Translate](#), and then configure CloudFront to accelerate the content generated by those services. For more information about using the plugin to set up Amazon Polly, see [WordPress Plugin for Amazon Polly](#) in the *Amazon*

Polly Developer Guide. For more information about using CloudFront to accelerate the content generated by Amazon Polly, see [\(Optional\) Create a CloudFront distribution for Amazon Polly content \(p. 26\)](#).

Topics

- [Prerequisites \(p. 20\)](#)
- [Step 1: Install the plugin \(p. 22\)](#)
- [Step 2: Configure and use CloudFront with the plugin \(p. 22\)](#)
- [\(Optional\) Deactivate site acceleration \(p. 24\)](#)
- [\(Optional\) Remove site acceleration and delete the CloudFront distribution \(p. 25\)](#)
- [\(Optional\) Deactivate and remove the plugin \(p. 25\)](#)
- [\(Optional\) Create a CloudFront distribution for Amazon Polly content \(p. 26\)](#)
- [Troubleshooting \(p. 26\)](#)

Prerequisites

To use the AWS for WordPress plugin, you need an AWS account, an AWS Identity and Access Management (IAM) user, and a WordPress website.

Topics

- [Creating an AWS account \(p. 20\)](#)
- [Creating an IAM user \(p. 20\)](#)
- [Creating a WordPress website \(p. 21\)](#)

Creating an AWS account

If you have an AWS account already, you can skip this section. Otherwise, create one.

To create 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.

Creating an IAM user

To use the AWS for WordPress plugin, you must create an *IAM user* for the plugin. An IAM user is a person or application under an AWS account that has permission to make API calls to AWS services.

Note

If you don't use [WordPress.com](#) and instead have a self-hosted WordPress website on Amazon EC2, you can use an IAM role instead of an IAM user. For more information, see [IAM Roles for Amazon EC2](#) in the *Amazon EC2 User Guide*.

The following procedure includes the steps to attach an IAM *policy* to the IAM user. An IAM policy is a document that defines the permissions that apply to the user.

To create an IAM user

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

2. In the navigation pane, choose **Users**. Then choose **Add user**.
3. On the **Set user details** page, do the following:
 - a. For **User name**, enter **AWSForWordPressPlugin**.
 - b. For **Access type**, choose **Programmatic access**.
 - c. Choose **Next: Permissions**.
4. On the **Set permissions** page, do the following:
 - a. Choose **Attach existing policies directly**.
 - b. In the search box, enter **WordPress**, and then select the check box next to **AWSForWordPressPluginPolicy**.
5. Choose **Next: Review**.
6. Choose **Create user**.
7. Choose **Download .csv** to save the user's credentials (access key ID and secret access key) to your computer. You need them to configure the AWS for WordPress plugin.

Note

The *AWSForWordPressPluginPolicy* is an AWS managed policy that gives the user permission to use all the features included in the AWS for WordPress plugin. When new features are added to the plugin, AWS will update this policy to include the permissions necessary to use the new features.

Important

This is the only time that you can save the user's secret access key, so make sure to save it now.

Protect the IAM user's credentials

The IAM user that you created in the preceding section can do the following in your AWS account:

- Create, modify, tag, list, and delete CloudFront distributions with the tag "createdBy" : "AWSForWordPressPlugin", and create and list invalidations in those distributions.
- Request, tag, list, and delete AWS Certificate Manager certificates in the US East (N. Virginia) Region.
- Create AWS CloudFormation stacks in the US East (N. Virginia) Region, and modify, list, and delete stacks with the tag "createdBy" : "AWSForWordPressPlugin".
- Use Amazon Polly to convert text into speech, and list the Amazon Polly voices that are available.
- Use Amazon Translate to translate text from one language to another.
- Determine whether a particular Amazon S3 bucket exists.
- Create Amazon S3 buckets whose names begins with `audio_for_wordpress` or `audio-for-wordpress`, and create, delete, and list objects in those buckets.

Important

To prevent unauthorized users from gaining these permissions, protect the IAM user's credentials. Treat the secret access key like a password; store it in a safe place, and don't share it with anyone. Like a password, [rotate the access key periodically](#). If the secret access key is accidentally leaked, [delete it immediately](#). Then you can create a new access key to use with the AWS for WordPress plugin.

Creating a WordPress website

If you have a WordPress website already, you can skip ahead to [Step 1: Install the plugin \(p. 22\)](#).

If you don't have a WordPress website, you can create one using WordPress.com. To use the AWS for WordPress plugin, you need a WordPress.com Business or eCommerce plan.

You can also install the WordPress software on your own web server, using Amazon Lightsail, Amazon EC2, or another web hosting platform. Hosting your own WordPress website involves more steps than using WordPress.com, and requires the ability to configure and manage a web server, a load balancer, DNS records, and web server certificates.

Regardless of how you set up your WordPress website, you need the following before you can use the AWS for WordPress plugin:

- Your website must have its own *domain name*. A domain name, also known as a *web address* or a *URL (uniform resource locator)*, is the address that visitors use to go to your website. For example, Amazon's domain name is *amazon.com*. In this topic, we use *example.com* as a generic example domain name, but you need a custom domain name for your website.
- Your website must work using HTTPS. This is a security best practice, and the plugin assumes that your website works using HTTPS. To check, go to your website's address using HTTPS (for example, ***https://example.com***) and make sure that your website displays correctly.

For a step-by-step tutorial that explains how to create a WordPress website on AWS using Amazon Lightsail, see [Accelerating WordPress with CloudFront using the AWS for WordPress Plugin](#) on the AWS Networking and Content Delivery blog.

When your website has a domain name and works using HTTPS, proceed to the following section.

Step 1: Install the plugin

Before you install the plugin, make sure to complete the [prerequisites \(p. 20\)](#).

To install the plugin

1. Log in to the admin dashboard for your WordPress website, also known as **WP Admin**.
2. Choose **Plugins**.
3.
 - If you already have the **Amazon AI** (the plugin's previous name) or **AWS for WordPress** plugin:
 1. Select the check box next to **Amazon AI** or **AWS for WordPress**.
 2. In the **Bulk Action** menu, choose **Update**, and then choose **Apply**.
 - If you don't have the **Amazon AI** or **AWS for WordPress** plugin:
 1. Choose **Add New**.
 2. In the search box, enter **AWS for WordPress**.
 3. Find the **AWS for WordPress** plugin. Choose **Install Now**, and then choose **Activate**.

After you activate the plugin, proceed to the following section to configure and use it.


Step 2: Configure and use CloudFront with the plugin

When you use CloudFront with the AWS for WordPress plugin for site acceleration, the plugin uses a *subdomain*, also known as an *alternate domain name* or *CNAME*, to send your website's traffic through CloudFront. This can reduce latency and improve the viewing experience by loading resources faster.

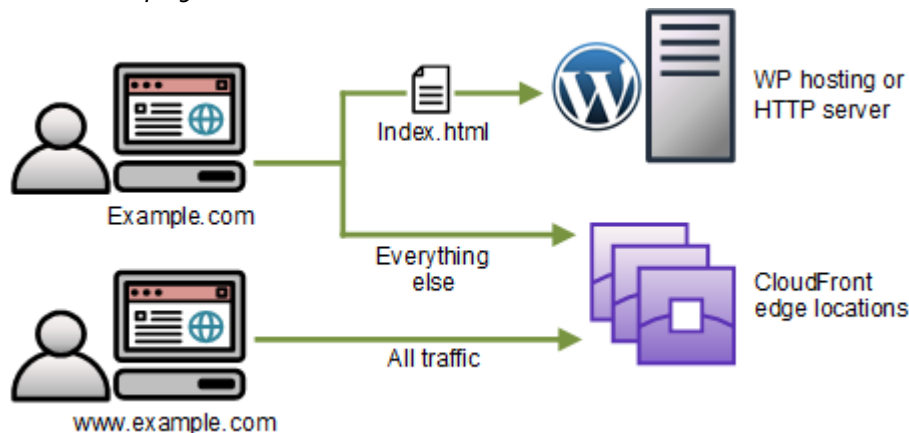
Without the plugin's site acceleration, all the traffic of your website's viewers goes to the server that hosts your WordPress website. After completing the steps in the following procedure, you can enable the plugin's site acceleration, which gives viewers two options for visiting your website:

- When viewers use your website's domain name, such as *example.com*, all the traffic goes through CloudFront, except for the website's index page and a few small image files.
- When viewers use your website's alternate domain name, such as *www.example.com*, all the traffic goes through CloudFront.

Using either domain, your website's viewers get lower latency and a faster, more reliable viewing experience. We recommend telling viewers to use your website's alternate domain name. The following diagrams show your viewers' traffic with and without the plugin's site acceleration.



Without the plugin's site acceleration



With the plugin's site acceleration

To configure and use CloudFront with the plugin (enable site acceleration)

1. Log in to the admin dashboard for your WordPress website, also known as **WP Admin**.
2. In the left navigation, choose **AWS**.
3. Paste or enter the access key ID and secret access key [that you saved previously \(p. 20\)](#), and then choose **Save Changes**.

Note

If you host WordPress on Amazon EC2, you can skip this step and use an IAM role instead of an IAM user. In that case, keep these two fields blank. For more information about IAM roles, see [IAM Roles for Amazon EC2](#) in the *Amazon EC2 User Guide*.

Note

Regardless of what you choose for **AWS Region**, the plugin's CloudFront feature creates all resources in the US East (N. Virginia) Region.

4. In the navigation pane, choose **CloudFront**.
5. On the **CloudFront Setup** page, do the following:
 - a. If necessary, for **Origin Domain Name**, enter your website's domain name, for example, *example.com*.

- b. For **CloudFront Alternate Domain Name**, enter a subdomain that viewers will use for your website's accelerated experience. We recommend using `www` in front of your website's domain name, for example, `www.example.com`.
 - c. Choose **Initiate Setup**.
6. CloudFront uses AWS Certificate Manager to create a certificate for your alternate domain name, and you must validate the certificate within 72 hours of the request. Do this by adding the DNS record that the plugin shows on the setup page. The process for adding this validation record varies depending on your DNS service provider. If you use WordPress.com hosting services, see [their documentation](#) for information about how to update DNS records with a custom entry. If you use Amazon Route 53 for DNS, see [Creating Records by Using the Amazon Route 53 Console](#) in the *Amazon Route 53 Developer Guide*.

After you add the DNS record, return to the setup page and choose **Check status of SSL certificate**.

When you complete this step, CloudFront sets up a distribution that is optimized for WordPress. This process can take some time to deploy globally. The setup page automatically refreshes every ten seconds to keep you updated while the deployment is in progress.

7. After the deployment is complete, create a DNS record to point your alternative domain name (for example, `www.example.com`) to your new CloudFront distribution. Do this by adding the DNS record that the plugin shows on the setup page. The process for adding this validation record varies depending on your DNS service provider. If you use WordPress.com hosting services, see [their documentation](#) for information about how to update DNS records with a custom entry. If you use Amazon Route 53 for DNS, see [Creating Records by Using the Amazon Route 53 Console](#) in the *Amazon Route 53 Developer Guide*.

After you add the DNS record, return to the setup page and choose **Check status of CloudFront DNS record**.

8. Choose **Activate Site Acceleration**, and then choose **Save Changes**.

When you activate site acceleration, the AWS for WordPress plugin configures your website to serve the website's resources—for example, CSS and JavaScript files, and images—from your CloudFront distribution. You can verify that the plugin accelerates your website for viewers by viewing your website from a private browsing window, or by using a different browser outside of WordPress's admin mode. Make sure that you navigate to your website using the alternate domain name, for example, `www.example.com`.

(Optional) Deactivate site acceleration

You can deactivate site acceleration to serve all of your website's resources from your web server host, bypassing the CloudFront distribution. This leaves your distribution intact and available for use when you choose to reactivate the plugin's site acceleration.

Warning

Before you deactivate site acceleration, edit the DNS record for your alternative domain name (such as `www.example.com`) so that it points to your website's domain (such as `example.com`). If you don't do this first, you might experience downtime or problems with your website. After you edit the DNS record, wait longer than the record's time to live (TTL) value before deactivating site acceleration.

If you use WordPress.com hosting services, see [their documentation](#) for information about how to edit DNS records. If you use Amazon Route 53 for DNS, see [Editing Records](#) in the *Amazon Route 53 Developer Guide*.

To deactivate site acceleration

1. Log in to your WordPress website, and then choose **WP Admin**.
2. In the navigation pane, choose **AWS**.

3. In the navigation pane, choose **CloudFront**.
4. Clear the **Activate Site Acceleration** check box, and then choose **Save Changes**.

Deactivating site acceleration is reversible. To reactivate it, select the **Activate Site Acceleration** check box, and then choose **Save Changes**.

(Optional) Remove site acceleration and delete the CloudFront distribution

You can use the AWS for WordPress plugin to delete your CloudFront distribution. This is not reversible. To use the AWS for WordPress plugin for site acceleration again, you must [reconfigure the plugin \(p. 22\)](#), which creates a new CloudFront distribution.

Warning

Before you delete your CloudFront distribution, edit the DNS record for your alternative domain name (such as *www.example.com*) so that it points to your website's domain (such as *example.com*). If you don't do this first, you might experience downtime or problems with your website. After you edit the DNS record, wait longer than the record's time to live (TTL) value before deleting your CloudFront distribution.

If you use WordPress.com hosting services, see [their documentation](#) for information about how to edit DNS records. If you use Amazon Route 53 for DNS, see [Editing Records](#) in the *Amazon Route 53 Developer Guide*.

To remove site acceleration and delete your CloudFront distribution

1. Log in to your WordPress website, and then choose **WP Admin**.
2. In the navigation pane, choose **AWS**.
3. In the navigation pane, choose **CloudFront**.
4. Choose **Remove Site Acceleration**, and then choose **OK**.

When you complete these steps, the AWS for WordPress plugin deletes your CloudFront distribution. This can take several minutes to complete. After the process is complete, you can optionally open the [AWS Management Console](#) to verify that the CloudFront, AWS Certificate Manager, and AWS CloudFormation resources created by the plugin are deleted.

(Optional) Deactivate and remove the plugin

You can deactivate the AWS for WordPress plugin to stop using all of its features for CloudFront and other AWS services. You can also delete the plugin to remove it from your WordPress website completely.

Warning

Before you deactivate and delete the plugin, edit the DNS record for your alternative domain name (such as *www.example.com*) so that it points to your website's domain (such as *example.com*). If you don't do this first, you might experience downtime or problems with your website. After you edit the DNS record, wait longer than the record's time to live (TTL) value before deactivating and deleting the plugin.

If you use WordPress.com hosting services, see [their documentation](#) for information about how to edit DNS records. If you use Amazon Route 53 for DNS, see [Editing Records](#) in the *Amazon Route 53 Developer Guide*.

Note

If you deactivate and delete the plugin without first removing site acceleration, the plugin does not delete the CloudFront, AWS Certificate Manager, and AWS CloudFormation resources that it created. These resources remain in your AWS account, and you are charged for any usage that

exceeds the [AWS Free Tier](#). To delete these resources before deleting the plugin, see [\(Optional\) Remove site acceleration and delete the CloudFront distribution \(p. 25\)](#).

To deactivate the AWS for WordPress plugin

1. Log in to your WordPress website, and then choose **WP Admin**.
2. Choose **Plugins**.
3. Locate the **AWS for WordPress** plugin, and then choose **Deactivate**.

Deactivating the plugin is reversible. To reactivate it, choose **Activate**.

4. To completely remove the AWS for WordPress plugin, choose **Delete**.

(Optional) Create a CloudFront distribution for Amazon Polly content

If you use the AWS for WordPress plugin with Amazon Polly, you can create a CloudFront distribution to accelerate the audio content generated by Amazon Polly. For more information about using the plugin with Amazon Polly, see [WordPress Plugin for Amazon Polly](#) in the *Amazon Polly Developer Guide*.

To create a CloudFront distribution for Amazon Polly audio

1. Log in to the admin dashboard for your WordPress website, also known as **WP Admin**.
2. In the left navigation, choose **AWS**.
3. In the **Cloud Storage** section, make note of your S3 bucket name. It will begin with **audio-for-wordpress** or **audio_for_wordpress**. You need this bucket name to complete the following steps.
4. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
5. Choose **Create Distribution**.
6. Choose **Get Started** for a **Web** distribution.
7. For **Origin Domain Name**, choose the Amazon S3 bucket whose name you noted in a previous step.
8. Scroll to the bottom of the page, and then choose **Create Distribution**.
9. Choose the distribution that you created in the previous step, and then make note of the distribution's **Domain Name**. You need this domain name to complete the following steps.
10. Log in to the admin dashboard for your WordPress website, also known as **WP Admin**.
11. In the left navigation, choose **AWS**.
12. For **Amazon CloudFront (CDN) domain name**, enter the domain name that you noted in a previous step.
13. Choose **Save Changes**.

Troubleshooting

If you encounter problems with the AWS for WordPress plugin, the following topics can help you solve them. To report bugs or to get help with other problems that are not covered by these topics, [open an issue on GitHub](#).

Topics

- [Can't connect to AWS \(p. 27\)](#)
- [User is not authorized \(p. 27\)](#)

- [CloudFront settings page is blank \(p. 27\)](#)
- [DescribeCertificate error \(p. 27\)](#)
- [AWS CloudFormation error \(p. 27\)](#)
- [CloudFront distribution deployment seems stuck \(p. 28\)](#)
- [Alternate domain is not working \(p. 28\)](#)
- [Can't find AWS resources \(p. 28\)](#)

Can't connect to AWS

The plugin might display the following error: Can't connect to AWS. Check your credentials and make sure your AWS account is active. If you see this error, try the following:

- Make sure that you entered your AWS access key and AWS secret key in the plugin's **General configuration** page. For more information, see [Step 2: Configure and use CloudFront with the plugin \(p. 22\)](#).
- Make sure that the IAM user that you created for the plugin has the correct permissions. For more information, see [Creating an IAM user \(p. 20\)](#).

User is not authorized

The plugin might display the following error messages:

- Error in Setup
- AccessDenied
- User: `<user ARN>` is not authorized to perform `<action>`

If you see one of these errors, make sure that the IAM user that you created for the plugin has the correct permissions. For more information, see [Creating an IAM user \(p. 20\)](#).

CloudFront settings page is blank

When you navigate to the plugin's CloudFront settings page, the page might be blank. This means that you haven't entered your AWS access key and AWS secret key in the plugin's **General configuration** page. For more information, see [Step 2: Configure and use CloudFront with the plugin \(p. 22\)](#).

DescribeCertificate error

The plugin might display the following error messages:

- Error in Setup
- Found 1 error while validating the input provided for the DescribeCertificate operation:
[CertificateArn] expected string length to be >= 20, but found string length of 0

If you see one of these errors, choose **Restart Setup**, and then make sure that you enter a domain name, not an IP address, for **Origin Domain Name** and **CloudFront Alternate Domain Name**. For more information, see [Step 2: Configure and use CloudFront with the plugin \(p. 22\)](#).

AWS CloudFormation error

The plugin might display the following error messages:

- Caught exception in method AmazonAl_Cloudformation
- Stack is in an unexpected state. CloudFront Distribution state is: `<distribution state>` and Stack state is: `stack state`

If you see one of these errors, choose **Restart Setup** to try again. If you're comfortable diagnosing errors using the [AWS CloudFormation console](#), you can open the console to see what went wrong.

CloudFront distribution deployment seems stuck

When setting up site acceleration, the plugin might show the **CloudFront Distribution Deployment** step for a long time, and it might seem like the plugin is stuck at this step. This step can take several minutes to complete. The plugin refreshes every ten seconds during this step, and displays a message like this: Last updated at `<date and time of last update>`. Look for this message to see when the plugin last refreshed the page. If it was within the last minute, we recommend that you continue to wait for this step to complete. If the plugin has not refreshed in a while, you can try reloading the page.

Alternate domain is not working

If you finished setting up CloudFront with the plugin but your alternate domain name (for example, *www.example.com*) isn't working, make sure that you added a CNAME record to your DNS records. If you use WordPress.com hosting services, see [their documentation](#) for information about how to update DNS records with a custom entry. If you use Amazon Route 53 for DNS, see [Creating Records by Using the Amazon Route 53 Console](#) in the *Amazon Route 53 Developer Guide*.

If you used the plugin to delete your CloudFront distribution and your alternate domain name (for example, *www.example.com*) isn't working, make sure that you updated your DNS records to repoint the alternate domain name to your website's apex domain (for example, *example.com*) and that you've waited longer than the DNS record's time to live (TTL) value. If you use WordPress.com hosting services, see [their documentation](#) for information about how to update DNS records with a custom entry. If you use Amazon Route 53 for DNS, see [Creating Records by Using the Amazon Route 53 Console](#) in the *Amazon Route 53 Developer Guide*.

Can't find AWS resources

The CloudFront feature of the plugin creates resources in several AWS services, including CloudFront, AWS Certificate Manager, and AWS CloudFormation. If you're looking for these resources in the [AWS Management Console](#) or listing them using an API, make sure that you use the US East (N. Virginia) Region (us-east-1).

Getting started with a secure static website

You can get started with Amazon CloudFront by using the solution described in this topic to create a secure static website for your domain name. A *static website* uses only static files—like HTML, CSS, JavaScript, images, and videos—and doesn't need servers or server-side processing. With this solution, your website gets the following benefits:

- **Uses the durable storage of [Amazon Simple Storage Service \(Amazon S3\)](#)** – This solution creates an Amazon S3 bucket to host your static website's content. To update your website, just upload your new files to the S3 bucket.
- **Is sped up by the Amazon CloudFront content delivery network** – This solution creates a CloudFront distribution to serve your website to viewers with low latency. The distribution is configured with an [origin access identity \(p. 209\)](#) to make sure that the website is accessible only through CloudFront, not directly from S3.

- **Is secured by HTTPS and additional security headers** – This solution creates an SSL/TLS certificate in [AWS Certificate Manager \(ACM\)](#), and attaches it to the CloudFront distribution. This certificate enables the distribution to serve your domain's website securely with HTTPS.

This solution also uses [Lambda@Edge](#) (p. 318) to add security headers to every server response. Security headers are a group of headers in the web server response that tell web browsers to take extra security precautions. For more information, refer to this blog post: [Adding HTTP Security Headers Using Lambda@Edge and Amazon CloudFront](#).

- **Is configured and deployed with AWS CloudFormation** – This solution uses an AWS CloudFormation template to set up all the components, so you can focus more on your website's content and less on configuring components.

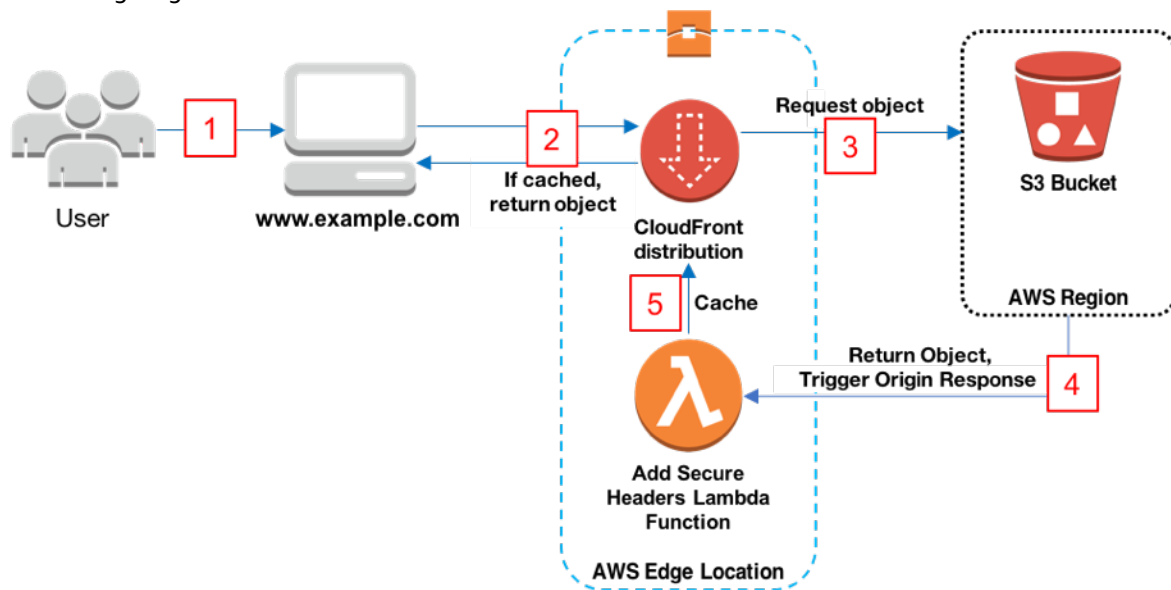
This solution is open source on GitHub. To view the code, submit a pull request, or open an issue, go to <https://github.com/aws-samples/amazon-cloudfront-secure-static-site>.

Topics

- [Solution overview](#) (p. 29)
- [Deploying the solution](#) (p. 30)

Solution overview

The following diagram shows an overview of how this static website solution works:



1. The viewer requests the website at www.example.com.
2. If the requested object is cached, CloudFront returns the object from its cache to the viewer.
3. If the object is not in CloudFront's cache, CloudFront requests the object from the origin (an S3 bucket).
4. S3 returns the object to CloudFront, which triggers the Lambda@Edge [origin response event](#) (p. 339).
5. The object, including the security headers added by the Lambda@Edge function, is added to CloudFront's cache.
6. (Not shown) The objects is returned to the viewer. Subsequent requests for the object that come to the same CloudFront edge location are served from the CloudFront cache.

Deploying the solution

To deploy this secure static website solution, you can choose from either of the following options:

- Use the AWS CloudFormation console to deploy the solution with default content, then upload your website content to Amazon S3.
- Clone the solution to your computer to add your website content. Then, deploy the solution with the AWS Command Line Interface (AWS CLI).

Topics

- [Prerequisites \(p. 30\)](#)
- [Using the AWS CloudFormation console \(p. 30\)](#)
- [Cloning the solution locally \(p. 31\)](#)
- [Finding access logs \(p. 32\)](#)

Prerequisites

To use this solution, you must have the following prerequisites:

- A registered domain name, such as example.com, that's pointed to an Amazon Route 53 hosted zone. The hosted zone must be in the same AWS account where you deploy this solution. If you don't have a registered domain name, you can [register one with Route 53](#). If you have a registered domain name but it's not pointed to a Route 53 hosted zone, [configure Route 53 as your DNS service](#).
- AWS Identity and Access Management (IAM) permissions to launch CloudFormation templates that create IAM roles, and permissions to create all the AWS resources in the solution.

You are responsible for the costs incurred while using this solution. For more information about costs, see [the pricing pages for each AWS service](#).

Using the AWS CloudFormation console

To deploy using the CloudFormation console

1. Choose **Launch on AWS** to open this solution in the AWS CloudFormation console. If necessary, sign in to your AWS account.



2. The **Create stack** wizard opens in the AWS CloudFormation console, with prepopulated fields that specify this solution's CloudFormation template.

At the bottom of the page, choose **Next**.

3. On the **Specify stack details** page, enter values for the following fields:
 - **SubDomain** – Enter the subdomain to use for your website. For example, if the subdomain is *www*, your website is available at *www.example.com*. (Replace example.com with your domain name, as explained in the following bullet.)
 - **DomainName** – Enter your domain name, such as *example.com*. This domain must be pointed to a Route 53 hosted zone.


When finished, choose **Next**.

4. (Optional) On the **Configure stack options** page, [add tags and other stack options](#).

When finished, choose **Next**.

5. On the **Review** page, scroll to the bottom of the page, then select the two boxes in the **Capabilities** section. These capabilities allow AWS CloudFormation to create an IAM role that allows access to the stack's resources, and to name the resources dynamically.
6. Choose **Create stack**.
7. Wait for the stack to finish creating. The stack creates some nested stacks, and can take several minutes to finish. When it's finished, the **Status** changes to **CREATE_COMPLETE**.

When the status is **CREATE_COMPLETE**, go to <https://www.example.com> to view your website (replace `www.example.com` with the subdomain and domain name that you specified in step 3). You should see the website's default content:



To replace the website's default content with your own

1. Open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. Choose the bucket whose name begins with **amazon-cloudfront-secure-static-site-s3bucketroot-**.

Note

Make sure to choose the bucket with **s3bucketroot** in its name, not **s3bucketlogs**. The bucket with **s3bucketroot** in its name contains the website content. The one with **s3bucketlogs** contains only log files.

3. Delete the website's default content, then upload your own.

Note

If you viewed your website with this solution's default content, then it's likely that some of the default content is cached in a CloudFront edge location. To make sure that viewers see your updated website content, *invalidate* the files to remove the cached copies from CloudFront edge locations. For more information, see [Invalidating Files \(p. 110\)](#).

Cloning the solution locally

Prerequisites

To add your website content before deploying this solution, you must package the solution's artifacts locally, which requires Node.js and npm. For more information, see <https://www.npmjs.com/get-npm>.

To add your website content and deploy the solution

1. Clone or download the solution from <https://github.com/aws-samples/amazon-cloudfront-secure-static-site>. After you clone or download it, open a command prompt or terminal and navigate to the `amazon-cloudfront-secure-static-site` folder.
2. Run the following command to install and package the solution's artifacts:

```
make package-function
```

3. Copy your website's content into the `www` folder, overwriting the default website content.

4. Run the following AWS CLI command to create an Amazon S3 bucket to store the solution's artifacts. Replace *example-bucket-for-artifacts* with your own bucket name.

```
aws s3 mb s3://example-bucket-for-artifacts --region us-east-1
```

5. Run the following AWS CLI command to package the solution's artifacts as an AWS CloudFormation template. Replace *example-bucket-for-artifacts* with the name of the bucket that you created in the previous step.

```
aws cloudformation package \  
  --region us-east-1 \  
  --template-file templates/main.yaml \  
  --s3-bucket example-bucket-for-artifacts \  
  --output-template-file packaged.template
```

6. Run the following command to deploy the solution with AWS CloudFormation, replacing the following values:

- *your-CloudFormation-stack-name* – Replace with a name for the AWS CloudFormation stack.
- *example.com* – Replace with your domain name. This domain must be pointed to a Route 53 hosted zone in the same AWS account.
- *www* – Replace with the subdomain to use for your website. For example, if the subdomain is *www*, your website is available at *www.example.com*.

```
aws cloudformation deploy \  
  --region us-east-1 \  
  --stack-name your-CloudFormation-stack-name \  
  --template-file packaged.template \  
  --capabilities CAPABILITY_NAMED_IAM CAPABILITY_AUTO_EXPAND \  
  --parameter-overrides DomainName=example.com SubDomain=www
```

7. Wait for the AWS CloudFormation stack to finish creating. The stack creates some nested stacks, and can take several minutes to finish. When it's finished, the **Status** changes to **CREATE_COMPLETE**.

When the status changes to **CREATE_COMPLETE**, go to <https://www.example.com> to view your website (replace *www.example.com* with the subdomain and domain name that you specified in the previous step). You should see your website's content.

Finding access logs

This solution enables [access logs \(p. 439\)](#) for the CloudFront distribution. Complete the following steps to locate the distribution's access logs.

To locate the distribution's access logs

1. Open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. Choose the bucket whose name begins with **amazon-cloudfront-secure-static-site-s3bucketlogs-**.

Note

Make sure to choose the bucket with **s3bucketlogs** in its name, not **s3bucketroot**. The bucket with **s3bucketlogs** in its name contains log files. The one with **s3bucketroot** contains the website content.

3. The folder named **cdn** contains the CloudFront access logs.

Working with distributions

You create a CloudFront distribution to tell CloudFront where you want content to be delivered from, and the details about how to track and manage content delivery. The following topics explain some basics about CloudFront distributions and provide detailed information about the settings you can choose to configure your distributions to meet your business needs.

Topics

- [Overview of distributions \(p. 33\)](#)
- [Creating, Updating, and Deleting Distributions \(p. 36\)](#)
- [Using Amazon S3 Origins, MediaPackage Channels, and Custom Origins for Web Distributions \(p. 66\)](#)
- [Using Custom URLs for Files by Adding Alternate Domain Names \(CNAMEs\) \(p. 71\)](#)
- [Using WebSocket with CloudFront Distributions \(p. 81\)](#)

Overview of distributions

When you want to use CloudFront to distribute your content, you create a distribution and choose the configuration settings you want. For example:

- **Your content origin**—that is, the Amazon S3 bucket, MediaPackage channel, or HTTP server from which CloudFront gets the files to distribute. You can specify any combination of up to 25 Amazon S3 buckets, channels, and/or HTTP servers as your origins.
- **Access**—whether you want the files to be available to everyone or restrict access to some users.
- **Security**—whether you want CloudFront to require users to use HTTPS to access your content.
- **Cache key**—which values, if any, you want to include in the *cache key*. The cache key uniquely identifies each file in the cache for a given distribution.
- **Origin request settings**—whether you want CloudFront to include HTTP headers, cookies, or query strings in requests that it sends to your origin.
- **Geo-restrictions**—whether you want CloudFront to prevent users in selected countries from accessing your content.
- **Access logs**—whether you want CloudFront to create access logs that show viewer activity.

For the current maximum number of distributions that you can create for each AWS account, see [General Quotas on Web Distributions \(p. 497\)](#) and [Quotas on RTMP Distributions \(p. 502\)](#). There is no maximum number of files that you can serve per distribution.

You can use distributions to serve the following content over HTTP or HTTPS:

- Static and dynamic download content, for example, .html, .css, .js, and image files, using HTTP or HTTPS.
- Video on demand in different formats, such as Apple HTTP Live Streaming (HLS) and Microsoft Smooth Streaming. For more information, see the [Delivering Video on Demand \(VOD\) with CloudFront \(p. 298\)](#).

You can't serve Adobe Flash multimedia content over HTTP or HTTPS, but you can serve it using a CloudFront RTMP distribution. See [RTMP Distributions \(p. 305\)](#).

- A live event, such as a meeting, conference, or concert, in real time. For live streaming, you can create the distribution automatically by using an AWS CloudFormation stack. For more information, see [Delivering Live Streaming Video with CloudFront and AWS Media Services](#) (p. 300).

For information about creating a distribution, see [Steps for Creating a Distribution \(Overview\)](#) (p. 36).

Actions you can use with distributions

The following table lists the CloudFront actions that you can take to work with distributions and provides links to the corresponding documentation on how to do the actions with the CloudFront console and the CloudFront API.

Action	Using the CloudFront console	Using the CloudFront API: web distributions	Using the CloudFront API: RTMP distributions
Create a distribution	Web distributions: See Steps for Creating a Distribution (Overview) (p. 36) RTMP distributions: See Task List for Streaming Media Files Using RTMP (p. 307)	Go to CreateDistribution	Go to CreateStreamingDistribution
List your distributions	See Updating a Distribution (p. 63)	Go to ListDistributions	Go to ListStreamingDistributions
Get all information about a distribution	See Updating a Distribution (p. 63)	Go to GetDistribution	Go to GetStreamingDistribution
Get the distribution configuration	See Updating a Distribution (p. 63)	Go to GetDistributionConfig	Go to GetStreamingDistributionConfig
Update a distribution	See Updating a Distribution (p. 63)	Go to UpdateDistribution	Go to UpdateStreamingDistribution
Delete a distribution	See Deleting a Distribution (p. 65)	Go to DeleteDistribution	Go to DeleteStreamingDistribution

Required fields for creating and updating distributions

When you update a distribution by using the [UpdateDistribution](#) CloudFront API action, there are more required fields than when you create a distribution by using [CreateDistribution](#). Review the following tables for a summary of the fields that are required for creating and for updating a distribution.

DistributionConfig

Members	Required in Create API Call	Required in Update API Call
CallerReference	Y	Y

Members	Required in Create API Call	Required in Update API Call
Aliases	-	Y
DefaultRootObject	-	Y
Origins	Y	Y
DefaultCacheBehavior	Y	Y
CacheBehaviors	-	Y
CustomErrorResponses	-	Y
Comment	Y	Y
Logging	-	Y
PriceClass	-	Y
Enabled	Y	Y
ViewerCertificate	-	Y
Restrictions	-	Y
WebACLId	-	Y
HttpVersion	-	Y
IsIPV6Enabled	-	-

CacheBehavior

Members	Required in Create API Call	Required in Update API Call
PathPattern	Y	Y
TargetOriginId	Y	Y
ForwardedValues	Y	Y
TrustedSigners	Y	Y
ViewerProtocolPolicy	Y	Y
MinTTL	Y	Y
AllowedMethods	-	Y
SmoothStreaming	-	Y
DefaultTTL	-	Y
MaxTTL	Y	Y
Compress	-	Y
LambdaFunctionAssociations	-	Y
FieldLevelEncryptionId	-	Y

Creating, Updating, and Deleting Distributions

You can create, update, or delete a distribution by completing the steps in the following topics.

Topics

- [Steps for Creating a Distribution \(Overview\) \(p. 36\)](#)
- [Creating a Distribution \(p. 37\)](#)
- [Values That You Specify When You Create or Update a Distribution \(p. 38\)](#)
- [Values That CloudFront Displays in the Console \(p. 61\)](#)
- [Testing a Distribution \(p. 62\)](#)
- [Updating a Distribution \(p. 63\)](#)
- [Tagging Amazon CloudFront Distributions \(p. 64\)](#)
- [Deleting a Distribution \(p. 65\)](#)

Steps for Creating a Distribution (Overview)

The following task list summarizes the process for creating a distribution.

To Create a Distribution

1. Create one or more Amazon S3 buckets or configure HTTP servers as your origin servers. An origin is the location where you store the original version of your content. When CloudFront gets a request for your files, it goes to the origin to get the files that it distributes at edge locations. You can use any combination of Amazon S3 buckets and HTTP servers as your origin servers.

If you're using Amazon S3, note that the name of your bucket must be all lowercase and cannot contain spaces.

If you're using an Amazon EC2 server or another custom origin, review [Using Amazon EC2 or Other Custom Origins \(p. 68\)](#).

For the current maximum number of origins that you can create for a distribution, or to request a higher quota (formerly known as limit), see [General Quotas on Web Distributions \(p. 497\)](#).

2. Upload your content to your origin servers. If you don't want to restrict access to your content using CloudFront signed URLs, make the objects publicly readable.

Important

You are responsible for ensuring the security of your origin server. You must ensure that CloudFront has permission to access the server and that the security settings are appropriate to safeguard your content.

3. Create your CloudFront distribution:
 - For more information about creating a distribution using the CloudFront console, see [Creating a Distribution \(p. 37\)](#).
 - For information about creating a distribution using the CloudFront API, go to [CreateDistribution](#) in the *Amazon CloudFront API Reference*.
4. Optional: If you created your distribution using the CloudFront console, create more cache behaviors or origins for your distribution. For more information, see [To Update a CloudFront Distribution \(p. 63\)](#).
5. Test your distribution. For more information, see [Testing a Distribution \(p. 62\)](#).
6. Develop your website or application to access your content using the domain name that CloudFront returned after you created your distribution in Step 3. For example, if CloudFront returns

d111111abcdef8.cloudfront.net as the domain name for your distribution, the URL for the file `image.jpg` in an Amazon S3 bucket or in the root directory on an HTTP server will be `http://d111111abcdef8.cloudfront.net/image.jpg`.

If you specified one or more alternate domain names (CNAMEs) when you created your distribution, you can use your own domain name. In that case, the URL for `image.jpg` might be `http://www.example.com/image.jpg`.

Note the following:

- If you want to use signed URLs to restrict access to your content, see [Serving Private Content with Signed URLs and Signed Cookies \(p. 145\)](#).
- If you want to serve compressed content, see [Serving compressed files \(p. 116\)](#).
- For information about CloudFront request and response behavior for Amazon S3 and custom origins, see [Request and Response Behavior \(p. 263\)](#).

Creating a Distribution

You can create or update a distribution by using the CloudFront console or programmatically. This topic is about working with distributions by using the console.

If you want to create or update a distribution by using the CloudFront API, see [Create Distribution](#) or [Update Distribution](#) in the *Amazon CloudFront API Reference*.

Important

When you update your distribution, be aware that a number of additional fields are required that are not required to create a distribution. To help make sure that all of the required fields are included when you update your distribution by using the CloudFront API, follow the steps described in the *Amazon CloudFront API Reference*.

To see the current maximum number of distributions that you can create for each AWS account, or to request a higher quota (formerly known as limit), see [General Quotas on Web Distributions \(p. 497\)](#).

To create a CloudFront web distribution (console)

1. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
2. Choose **Create Distribution**.
3. On the first page of the **Create Distribution Wizard**, in the **Web** section, choose **Get Started**.
4. Specify settings for the distribution. For more information, see [Values That You Specify When You Create or Update a Distribution \(p. 38\)](#).
5. Save changes.
6. After CloudFront creates your distribution, the value of the **Status** column for your distribution will change from **InProgress** to **Deployed**. If you chose to enable the distribution, it will be ready to process requests after the status switches to **Deployed**.

The domain name that CloudFront assigns to your distribution appears in the list of distributions. (It also appears on the **General** tab for a selected distribution.)

Tip

You can use an alternate domain name, instead of the name assigned to you by CloudFront; by following the steps in [Using Custom URLs for Files by Adding Alternate Domain Names \(CNAMEs\) \(p. 71\)](#).

7. When your distribution is deployed, confirm that you can access your content using your new CloudFront URL or CNAME. For more information, see [Testing a Distribution \(p. 62\)](#).

To update a distribution (for example, to add or change cache behaviors), see [Updating a Distribution \(p. 63\)](#).

Values That You Specify When You Create or Update a Distribution

When you use the [CloudFront console](#) to create a new distribution or update an existing distribution, you specify the following values.

For more information about creating or updating a distribution by using the CloudFront console, see [Creating a Distribution \(p. 37\)](#) or [Updating a Distribution \(p. 63\)](#).

Delivery method

You specify the delivery method when you create a distribution. You can't change the delivery method for an existing distribution. You can choose **Web** or **RTMP**. Choose **Web**, unless you're using Adobe Flash Media Server with RTMP.

Note

[Adobe designated Flash as end-of-life at the end of 2020](#). As a result, Amazon CloudFront will no longer support Adobe Flash Media Server and will be deprecating Real-Time Messaging Protocol (RTMP) distributions by December 31, 2020. For more information, [read the full announcement on the Amazon CloudFront discussion forum](#).

Origin Settings (p. 40)

The following values apply to all types of origins:

- [Origin Domain Name \(p. 40\)](#)
- [Origin Path \(p. 42\)](#)
- [Origin ID \(p. 42\)](#)
- [Origin Connection Attempts \(p. 42\)](#)
- [Origin Connection Timeout \(p. 43\)](#)
- [Origin Custom Headers \(p. 43\)](#)

The following values apply only to Amazon S3 origins (those that are *not* using the S3 static website endpoint):

- [Restrict Bucket Access \(p. 43\)](#)
- [Origin Access Identity \(p. 44\)](#) (Applies only when you choose **Yes** for **Restrict Bucket Access**)
- [Comment \(p. 44\)](#) (Applies only when you choose **Create a New Identity** for **Origin Access Identity**)
- [Your Identities \(p. 44\)](#) (Applies only when you choose **Use an Existing Identity** for **Origin Access Identities**)
- [Grant Read Permissions on Bucket \(p. 44\)](#) (Applies only when you choose **Yes** for **Restrict Bucket Access**)

The following values apply only to custom origins, such as Amazon EC2 instances, Elastic Load Balancing load balancers, MediaPackage origins, MediaStore containers, or your own web server:

- [Minimum Origin SSL Protocol \(p. 44\)](#)
- [Origin Protocol Policy \(p. 45\)](#)
- [Origin Response Timeout \(p. 45\)](#)

- [Origin Keep-alive Timeout \(p. 46\)](#)
- [HTTP Port \(p. 46\)](#)
- [HTTPS Port \(p. 46\)](#)

Cache Behavior Settings (p. 46)

The following values apply to the **Default Cache Behavior Settings** (when you create a distribution) and to other cache behaviors that you create later.

- [Path Pattern \(p. 47\)](#)
- [Origin or Origin Group \(p. 48\)](#) (Applies only when you create or update a cache behavior for an existing distribution)
- [Viewer Protocol Policy \(p. 49\)](#)
- [Allowed HTTP Methods \(p. 49\)](#)
- [Field Level Encryption Config \(p. 49\)](#)
- [Cached HTTP Methods \(p. 50\)](#)
- [Cache Based on Selected Request Headers \(p. 50\)](#)
- [Whitelist Headers \(p. 50\)](#) (Applies only when you choose **Whitelist** for **Cache Based on Selected Request Headers**)
- [Object Caching \(p. 50\)](#)
- [Minimum TTL \(p. 50\)](#)
- [Maximum TTL \(p. 51\)](#)
- [Default TTL \(p. 51\)](#)
- [Forward Cookies \(p. 51\)](#)
- [Whitelist Cookies \(p. 51\)](#) (Applies only when you choose **Whitelist** for **Forward Cookies**)
- [Query String Forwarding and Caching \(p. 52\)](#)
- [Query String Whitelist \(p. 52\)](#) (Applies only when you choose **Forward all, cache based on whitelist** for **Query String Forwarding and Caching**)
- [Smooth Streaming \(p. 52\)](#)
- [Restrict Viewer Access \(Use Signed URLs or Signed Cookies\) \(p. 53\)](#)
- [Trusted Signers \(p. 53\)](#) (Applies only when you choose **Yes** for **Restrict Viewer Access (Use Signed URLs or Signed Cookies)**)
- [AWS Account Numbers \(p. 53\)](#) (Applies only when you choose **Specify Accounts** for **Trusted Signers**)
- [Compress Objects Automatically \(p. 53\)](#)

The following values apply to **Lambda Function Associations**.

- [CloudFront Event \(p. 54\)](#)
- [Lambda Function ARN \(p. 54\)](#)
- [Include Body \(p. 367\)](#)

Distribution Settings (p. 54)

- [Price Class \(p. 54\)](#)
- [AWS WAF Web ACL \(p. 54\)](#)
- [Alternate Domain Names \(CNAMEs\) \(p. 54\)](#)
- [SSL Certificate \(p. 55\)](#)

- [Custom SSL Client Support \(p. 56\)](#) (Applies only when you choose **Custom SSL Certificate (example.com)** for **SSL Certificate**)
- [Security Policy \(p. 56\)](#) (Minimum SSL/TLS version)
- [Supported HTTP Versions \(p. 57\)](#)
- [Default Root Object \(p. 57\)](#)
- [Logging \(p. 58\)](#)
- [Bucket for Logs \(p. 58\)](#)
- [Log Prefix \(p. 58\)](#)
- [Cookie Logging \(p. 58\)](#)
- [Enable IPv6 \(p. 59\)](#)
- [Comment \(p. 59\)](#)
- [Distribution State \(p. 59\)](#)

Custom Error Pages and Error Caching (p. 60)

- [HTTP Error Code \(p. 60\)](#)
- [Response Page Path \(p. 60\)](#)
- [HTTP Response Code \(p. 60\)](#)
- [Error Caching Minimum TTL \(seconds\) \(p. 60\)](#)

Restrictions (p. 61)

- [Enable Geo-Restriction \(p. 61\)](#)
- [Restriction Type \(p. 61\)](#)
- [Countries \(p. 61\)](#)

Origin Settings

When you create or update a distribution, you provide information about one or more locations—known as origins—where you store the original versions of your web content. CloudFront gets your web content from your origins and serves it to viewers via a world-wide network of edge servers. Each origin is either an Amazon S3 bucket or an HTTP server, for example, a web server.

For the current maximum number of origins that you can create for a distribution, or to request a higher quota (formerly known as limit), see [General Quotas on Web Distributions \(p. 497\)](#).

If you want to delete an origin, you must first edit or delete the cache behaviors that are associated with that origin.

Important

If you delete an origin, confirm that files that were previously served by that origin are available in another origin and that your cache behaviors are now routing requests for those files to the new origin.

When you create or update a distribution, you specify the following values for each origin.

Origin Domain Name

The DNS domain name of the Amazon S3 bucket or HTTP server from which you want CloudFront to get objects for this origin, for example:

- **Amazon S3 bucket** – `awsexamplebucket.s3.us-west-2.amazonaws.com`

Note

If you recently created the S3 bucket, the CloudFront distribution might return HTTP 307 Temporary Redirect responses for up to 24 hours. It can take up to 24 hours for the S3 bucket name to propagate to all AWS Regions. When the propagation is complete, the distribution automatically stops sending these redirect responses; you don't need to take any action. For more information, see [Why am I getting an HTTP 307 Temporary Redirect response from Amazon S3?](#) and [Temporary Request Redirection](#).

- **Amazon S3 bucket configured as a website** – `https://awsexamplebucket.s3-website.us-west-2.amazonaws.com`
- **MediaStore container** – `examplemediastore.data.mediastore.us-west-1.amazonaws.com`
- **MediaPackage endpoint** – `examplemediapackage.mediapackage.us-west-1.amazonaws.com`
- **Amazon EC2 instance** – `ec2-203-0-113-25.compute-1.amazonaws.com`
- **Elastic Load Balancing load balancer** – `example-load-balancer-1234567890.us-west-2.elb.amazonaws.com`
- **Your own web server** – `https://example.com`

Choose the domain name in the **Origin Domain Name** field, or type the name. The domain name is not case-sensitive.

If your origin is an Amazon S3 bucket, note the following:

- If the bucket is configured as a website, enter the Amazon S3 static website hosting endpoint for your bucket; don't select the bucket name from the list in the **Origin Domain Name** field. The static website hosting endpoint appears in the Amazon S3 console, on the **Properties** page under **Static Website Hosting**. For more information, see [Using Amazon S3 Buckets Configured as Website Endpoints for Your Origin](#) (p. 68).
- If you configured Amazon S3 Transfer Acceleration for your bucket, do not specify the `s3-accelerate` endpoint for **Origin Domain Name**.
- If you're using a bucket from a different AWS account and if the bucket is not configured as a website, enter the name, using the following format:

`bucket-name.s3.region.amazonaws.com`

If your bucket is in the US Standard Region and you want Amazon S3 to route requests to a facility in northern Virginia, use the following format:

`bucket-name.s3.us-east-1.amazonaws.com`

- The files must be publicly readable unless you secure your content in Amazon S3 by using a CloudFront origin access identity. For more information, see [Restricting Access to Amazon S3 Content by Using an Origin Access Identity](#) (p. 209).

Important

If the origin is an Amazon S3 bucket, the bucket name must conform to DNS naming requirements. For more information, go to [Bucket Restrictions and Limitations](#) in the *Amazon Simple Storage Service Developer Guide*.

When you change the value of **Origin Domain Name** for an origin, CloudFront immediately begins replicating the change to CloudFront edge locations. Until the distribution configuration is updated in a given edge location, CloudFront continues to forward requests to the previous HTTP server or Amazon S3 bucket. As soon as the distribution configuration is updated in that edge location, CloudFront begins to forward requests to the new HTTP server or Amazon S3 bucket.

Changing the origin does not require CloudFront to repopulate edge caches with objects from the new origin. As long as the viewer requests in your application have not changed, CloudFront continues to

serve objects that are already in an edge cache until the TTL on each object expires or until seldom-requested objects are evicted.

Origin Path

If you want CloudFront to request your content from a directory in your AWS resource or your custom origin, enter the directory path, beginning with a slash (/). CloudFront appends the directory path to the value of **Origin Domain Name**, for example, **cf-origin.example.com/production/images**. Do not add a slash (/) at the end of the path.

For example, suppose you've specified the following values for your distribution:

- **Origin Domain Name** – An Amazon S3 bucket named **myawsbucket**
- **Origin Path** – **/production**
- **Alternate Domain Names (CNAMEs)** – **example.com**

When a user enters **example.com/index.html** in a browser, CloudFront sends a request to Amazon S3 for **myawsbucket/production/index.html**.

When a user enters **example.com/acme/index.html** in a browser, CloudFront sends a request to Amazon S3 for **myawsbucket/production/acme/index.html**.

Origin ID

A string that uniquely distinguishes this origin or origin group in this distribution. If you create cache behaviors in addition to the default cache behavior, you use the ID that you specify here to identify the origin or origin group that you want CloudFront to route a request to when the request matches the path pattern for that cache behavior.

For more information, see the following:

- **Origins that you can specify:** [Using CloudFront Origin Groups \(p. 69\)](#)
- **Creating origin groups:** [Creating an Origin Group \(p. 232\)](#)
- **Working with cache behaviors:** [Cache Behavior Settings \(p. 46\)](#)

Origin Connection Attempts

The number of times that CloudFront attempts to connect to the origin. You can specify 1, 2, or 3 as the number of attempts. The default number (if you don't specify otherwise) is 3.

Use this setting together with **Origin Connection Timeout** to specify how long CloudFront waits before attempting to connect to the secondary origin or returning an error response to the viewer. By default, CloudFront waits as long as 30 seconds (3 attempts of 10 seconds each) before attempting to connect to the secondary origin or returning an error response. You can reduce this time by specifying fewer attempts, a shorter connection timeout, or both.

If the specified number of connection attempts fail, CloudFront does one of the following:

- If the origin is part of an origin group, CloudFront attempts to connect to the secondary origin. If the specified number of connection attempts to the secondary origin fail, then CloudFront returns an error response to the viewer.
- If the origin is not part of an origin group, CloudFront returns an error response to the viewer.

For a custom origin (including an Amazon S3 bucket that's configured with static website hosting), this setting also specifies the number of times that CloudFront attempts to get a response from the origin. For more information, see [Origin Response Timeout \(p. 45\)](#).

Origin Connection Timeout

The number of seconds that CloudFront waits when trying to establish a connection to the origin. You can specify a number of seconds between 1 and 10 (inclusive). The default timeout (if you don't specify otherwise) is 10 seconds.

Use this setting together with **Origin Connection Attempts** to specify how long CloudFront waits before attempting to connect to the secondary origin or before returning an error response to the viewer. By default, CloudFront waits as long as 30 seconds (3 attempts of 10 seconds each) before attempting to connect to the secondary origin or returning an error response. You can reduce this time by specifying fewer attempts, a shorter connection timeout, or both.

If CloudFront doesn't establish a connection to the origin within the specified number of seconds, CloudFront does one of the following:

- If the specified number of **Origin Connection Attempts** is more than 1, CloudFront tries again to establish a connection. CloudFront tries up to 3 times, as determined by the value of **Origin Connection Attempts**.
- If all the connection attempts fail and the origin is part of an origin group, CloudFront attempts to connect to the secondary origin. If the specified number of connection attempts to the secondary origin fail, then CloudFront returns an error response to the viewer.
- If all the connection attempts fail and the origin is not part of an origin group, CloudFront returns an error response to the viewer.

Origin Custom Headers

If you want CloudFront to add custom headers whenever it sends a request to your origin, specify the following values:

Header Name

The name of a header that you want CloudFront to add to requests that it sends to your origin.

Value

The value for the header that you specified in the **Custom Header** field.

For more information, see [Adding Custom Headers to Origin Requests \(p. 282\)](#).

For the current maximum number of custom headers that you can forward to the origin, the maximum length of a custom header name and value, and the maximum total length of all header names and values, see [Quotas \(p. 496\)](#).

Restrict Bucket Access

Note

This applies only to Amazon S3 bucket origins (those that are *not* using the S3 static website endpoint).

Choose **Yes** if you want to require users to access objects in an Amazon S3 bucket by using only CloudFront URLs, not by using Amazon S3 URLs. Then specify additional values.

Choose **No** if you want users to be able to access objects by using either CloudFront URLs or Amazon S3 URLs.

For more information, see [Restricting Access to Amazon S3 Content by Using an Origin Access Identity \(p. 209\)](#).

For information about how to require users to access objects on a custom origin by using only CloudFront URLs, see [Restricting Access to Files on Custom Origins \(p. 148\)](#).

Origin Access Identity

Note

This applies only to Amazon S3 bucket origins (those that are *not* using the S3 static website endpoint).

If you chose **Yes** for **Restrict Bucket Access**, choose whether to create a new origin access identity or use an existing one that is associated with your AWS account. If you already have an origin access identity, we recommend that you reuse it to simplify maintenance. For more information about origin access identities, see [Restricting Access to Amazon S3 Content by Using an Origin Access Identity \(p. 209\)](#).

Comment

Note

This applies only to Amazon S3 bucket origins (those that are *not* using the S3 static website endpoint).

If you chose **Create a New Identity** for **Origin Access Identity**, enter a comment that identifies the new origin access identity. CloudFront creates the origin access identity when you create this distribution.

Your Identities

Note

This applies only to Amazon S3 bucket origins (those that are *not* using the S3 static website endpoint).

If you chose **Use an Existing Identity** for **Origin Access Identity**, choose the origin access identity that you want to use. You cannot use an origin access identity that is associated with another AWS account.

Grant Read Permissions on Bucket

Note

This applies only to Amazon S3 bucket origins (those that are *not* using the S3 static website endpoint).

If you want CloudFront to automatically grant the origin access identity the permission to read objects in your Amazon S3 bucket, choose **Yes, Update Bucket Policy**.

Important

If you choose **Yes, Update Bucket Policy**, CloudFront updates the bucket policy to grant the specified origin access identity the permission to read objects in your bucket. However, CloudFront does not remove existing permissions in the bucket policy or permissions on individual objects. If users currently have permission to access the objects in your bucket using Amazon S3 URLs, they will still have that permission after CloudFront updates your bucket policy. To view or change the existing bucket policy and the existing permissions on the objects in your bucket, use a method provided by Amazon S3. For more information, see [Granting the OAI Permission to Read Files in Your Amazon S3 Bucket \(p. 212\)](#).

If you want to update permissions manually, for example, if you want to update ACLs on your objects instead of updating bucket permissions, choose **No, I will Update Permissions**.

Minimum Origin SSL Protocol

Note

This applies only to custom origins.

Choose the minimum TLS/SSL protocol that CloudFront can use when it establishes an HTTPS connection to your origin. Lower TLS protocols are less secure, so we recommend that you choose the latest TLS protocol that your origin supports.

If you use the CloudFront API to set the TLS/SSL protocol for CloudFront to use, you cannot set a minimum protocol. Instead, you specify all of the TLS/SSL protocols that CloudFront can use with your origin. For more information, see [OriginSslProtocols](#) in the *Amazon CloudFront API Reference*.

Origin Protocol Policy

Note

This applies only to custom origins.

The protocol policy that you want CloudFront to use when fetching objects from your origin server.

Choose one of the following values:

- **HTTP Only:** CloudFront uses only HTTP to access the origin.

Important

HTTP Only is the default setting when the origin is an Amazon S3 static website hosting endpoint, because Amazon S3 doesn't support HTTPS connections for static website hosting endpoints. The CloudFront console does not support changing this setting for Amazon S3 static website hosting endpoints.

- **HTTPS Only:** CloudFront uses only HTTPS to access the origin.
- **Match Viewer:** CloudFront communicates with your origin using HTTP or HTTPS, depending on the protocol of the viewer request. CloudFront caches the object only once even if viewers make requests using both HTTP and HTTPS protocols.

Important

For HTTPS viewer requests that CloudFront forwards to this origin, one of the domain names in the SSL certificate on your origin server must match the domain name that you specify for **Origin Domain Name**. Otherwise, CloudFront responds to the viewer requests with an HTTP status code 502 (Bad Gateway) instead of returning the requested object. For more information, see [Requirements for Using SSL/TLS Certificates with CloudFront \(p. 134\)](#).

Origin Response Timeout

Note

This applies only to custom origins.

The origin response timeout, also known as the *origin read timeout* or *origin request timeout*, applies to both of the following values:

- How long (in seconds) CloudFront waits for a response after forwarding a request to the origin.
- How long (in seconds) CloudFront waits after receiving a packet of a response from the origin and before receiving the next packet.

The default timeout is 30 seconds. You can change the value to be from 1 to 60 seconds. If you need a timeout value outside that range, [create a case in the AWS Support Center](#).

Tip

If you want to increase the timeout value because viewers are experiencing HTTP 504 status code errors, consider exploring other ways to eliminate those errors before changing the timeout value. See the troubleshooting suggestions in [HTTP 504 Status Code \(Gateway Timeout\) \(p. 259\)](#).

CloudFront behavior depends on the HTTP method in the viewer request:

- **GET** and **HEAD** requests – If the origin doesn't respond or stops responding within the duration of the response timeout, CloudFront drops the connection. CloudFront tries again to connect according to the value of [Origin Connection Attempts](#) (p. 42).
- **DELETE**, **OPTIONS**, **PATCH**, **PUT**, and **POST** requests – If the origin doesn't respond for the duration of the read timeout, CloudFront drops the connection and doesn't try again to contact the origin. The client can resubmit the request if necessary.

Origin Keep-alive Timeout

Note

This applies only to custom origins.

How long (in seconds) CloudFront tries to maintain a connection to your custom origin after it gets the last packet of a response. Maintaining a persistent connection saves the time that is required to re-establish the TCP connection and perform another TLS handshake for subsequent requests. Increasing the keep-alive timeout helps improve the request-per-connection metric for distributions.

Note

For the **Origin Keep-alive Timeout** value to have an effect, your origin must be configured to allow persistent connections.

The default timeout is 5 seconds. You can change the value to a number from 1 to 60 seconds. If you need a keep-alive timeout longer than 60 seconds, [create a case in the AWS Support Center](#).

HTTP Port

Note

This applies only to custom origins.

Optional. The HTTP port that the custom origin listens on. Valid values include ports 80, 443, and 1024 to 65535. The default value is port 80.

Important

Port 80 is the default setting when the origin is an Amazon S3 static website hosting endpoint, because Amazon S3 only supports port 80 for static website hosting endpoints. The CloudFront console does not support changing this setting for Amazon S3 static website hosting endpoints.

HTTPS Port

Note

This applies only to custom origins.

Optional. The HTTPS port that the custom origin listens on. Valid values include ports 80, 443, and 1024 to 65535. The default value is port 443.

Cache Behavior Settings

A cache behavior lets you configure a variety of CloudFront functionality for a given URL path pattern for files on your website. For example, one cache behavior might apply to all `.jpg` files in the `images` directory on a web server that you're using as an origin server for CloudFront. The functionality that you can configure for each cache behavior includes:

- The path pattern.
- If you have configured multiple origins for your CloudFront distribution, which origin you want CloudFront to forward your requests to.
- Whether to forward query strings to your origin.
- Whether accessing the specified files requires signed URLs.
- Whether to require users to use HTTPS to access those files.

- The minimum amount of time that those files stay in the CloudFront cache regardless of the value of any `Cache-Control` headers that your origin adds to the files.

When you create a new distribution, you specify settings for the default cache behavior, which automatically forwards all requests to the origin that you specify when you create the distribution. After you create a distribution, you can create additional cache behaviors that define how CloudFront responds when it receives a request for objects that match a path pattern, for example, `*.jpg`. If you create additional cache behaviors, the default cache behavior is always the last to be processed. Other cache behaviors are processed in the order in which they're listed in the CloudFront console or, if you're using the CloudFront API, the order in which they're listed in the `DistributionConfig` element for the distribution. For more information, see [Path Pattern \(p. 47\)](#).

When you create a cache behavior, you specify the one origin from which you want CloudFront to get objects. As a result, if you want CloudFront to distribute objects from all of your origins, you must have at least as many cache behaviors (including the default cache behavior) as you have origins. For example, if you have two origins and only the default cache behavior, the default cache behavior causes CloudFront to get objects from one of the origins, but the other origin is never used.

For the current maximum number of cache behaviors that you can add to a distribution, or to request a higher quota (formerly known as limit), see [General Quotas on Web Distributions \(p. 497\)](#).

Path Pattern

A path pattern (for example, `images/*.jpg`) specifies which requests you want this cache behavior to apply to. When CloudFront receives an end-user request, the requested path is compared with path patterns in the order in which cache behaviors are listed in the distribution. The first match determines which cache behavior is applied to that request. For example, suppose you have three cache behaviors with the following three path patterns, in this order:

- `images/*.jpg`
- `images/*`
- `*.gif`

Note

You can optionally include a slash (/) at the beginning of the path pattern, for example, `/images/*.jpg`. CloudFront behavior is the same with or without the leading `/`.

A request for the file `images/sample.gif` doesn't satisfy the first path pattern, so the associated cache behaviors are not applied to the request. The file does satisfy the second path pattern, so the cache behaviors associated with the second path pattern are applied even though the request also matches the third path pattern.

Note

When you create a new distribution, the value of **Path Pattern** for the default cache behavior is set to `*` (all files) and cannot be changed. This value causes CloudFront to forward all requests for your objects to the origin that you specified in the [Origin Domain Name \(p. 40\)](#) field. If the request for an object does not match the path pattern for any of the other cache behaviors, CloudFront applies the behavior that you specify in the default cache behavior.

Important

Define path patterns and their sequence carefully or you may give users undesired access to your content. For example, suppose a request matches the path pattern for two cache behaviors. The first cache behavior does not require signed URLs and the second cache behavior does require signed URLs. Users are able to access the objects without using a signed URL because CloudFront processes the cache behavior associated with the first match.

If you're working with a `MediaPackage` channel, you must include specific path patterns for the cache behavior that you define for the endpoint type for your origin. For example, for a DASH endpoint, you

type `*.mpd` for **Path Pattern**. For more information and specific instructions, see [Serving Live Video Formatted with AWS Elemental MediaPackage \(p. 301\)](#).

The path you specify applies to requests for all files in the specified directory and in subdirectories below the specified directory. CloudFront does not consider query strings or cookies when evaluating the path pattern. For example, if an `images` directory contains `product1` and `product2` subdirectories, the path pattern `images/*.jpg` applies to requests for any `.jpg` file in the `images`, `images/product1`, and `images/product2` directories. If you want to apply a different cache behavior to the files in the `images/product1` directory than the files in the `images` and `images/product2` directories, create a separate cache behavior for `images/product1` and move that cache behavior to a position above (before) the cache behavior for the `images` directory.

You can use the following wildcard characters in your path pattern:

- `*` matches 0 or more characters.
- `?` matches exactly 1 character.

The following examples show how the wildcard characters work:

Path pattern	Files that match the path pattern
<code>*.jpg</code>	All <code>.jpg</code> files
<code>images/*.jpg</code>	All <code>.jpg</code> files in the <code>images</code> directory and in subdirectories under the <code>images</code> directory
<code>a*.jpg</code>	<ul style="list-style-type: none">• All <code>.jpg</code> files for which the file name begins with <code>a</code>, for example, <code>apple.jpg</code> and <code>appalachian_trail_2012_05_21.jpg</code>• All <code>.jpg</code> files for which the file path begins with <code>a</code>, for example, <code>abra/cadabra/magic.jpg</code>.
<code>a??.jpg</code>	All <code>.jpg</code> files for which the file name begins with <code>a</code> and is followed by exactly two other characters, for example, <code>ant.jpg</code> and <code>abe.jpg</code>
<code>*.doc*</code>	All files for which the file name extension begins with <code>.doc</code> , for example, <code>.doc</code> , <code>.docx</code> , and <code>.docm</code> files. You can't use the path pattern <code>*.doc?</code> in this case, because that path pattern wouldn't apply to requests for <code>.doc</code> files; the <code>?</code> wildcard character replaces exactly one character.

The maximum length of a path pattern is 255 characters. The value can contain any of the following characters:

- A-Z, a-z

Path patterns are case-sensitive, so the path pattern `*.jpg` doesn't apply to the file `LOGO.JPG`.

- 0-9
- `_ - . * $ / ~ ' ' @ : +`
- `&`, passed and returned as `&`

Origin or Origin Group

Enter the value of an existing origin or origin group. This identifies the origin or origin group that you want CloudFront to route requests to when a request (such as `https://example.com/logo.jpg`) matches the path pattern for a cache behavior (such as `*.jpg`) or for the default cache behavior (`*`).

Viewer Protocol Policy

Choose the protocol policy that you want viewers to use to access your content in CloudFront edge locations:

- **HTTP and HTTPS:** Viewers can use both protocols.
- **Redirect HTTP to HTTPS:** Viewers can use both protocols, but HTTP requests are automatically redirected to HTTPS requests.
- **HTTPS Only:** Viewers can only access your content if they're using HTTPS.

For more information, see [Requiring HTTPS for Communication Between Viewers and CloudFront \(p. 122\)](#).

Allowed HTTP Methods

Specify the HTTP methods that you want CloudFront to process and forward to your origin:

- **GET, HEAD:** You can use CloudFront only to get objects from your origin or to get object headers.
- **GET, HEAD, OPTIONS:** You can use CloudFront only to get objects from your origin, get object headers, or retrieve a list of the options that your origin server supports.
- **GET, HEAD, OPTIONS, PUT, POST, PATCH, DELETE:** You can use CloudFront to get, add, update, and delete objects, and to get object headers. In addition, you can perform other POST operations such as submitting data from a web form.

Note

CloudFront caches responses to **GET** and **HEAD** requests and, optionally, **OPTIONS** requests. CloudFront does not cache responses to requests that use the other methods.

If you use an Amazon S3 bucket as the origin for your distribution and if you use CloudFront origin access identities, **POST** requests aren't supported in some Amazon S3 Regions and **PUT** requests in those Regions require an additional header. For more information, see [Using an OAI in Amazon S3 Regions that Support Only Signature Version 4 Authentication \(p. 215\)](#).

Important

If you choose **GET, HEAD, OPTIONS** or **GET, HEAD, OPTIONS, PUT, POST, PATCH, DELETE**, you might need to restrict access to your Amazon S3 bucket or to your custom origin to prevent users from performing operations that you don't want them to perform. The following examples explain how to restrict access:

- **If you're using Amazon S3 as an origin for your distribution:** Create a CloudFront origin access identity to restrict access to your Amazon S3 content, and grant permissions to the origin access identity. For example, if you configure CloudFront to accept and forward these methods *only* because you want to use **PUT**, you must still configure Amazon S3 bucket policies or ACLs to handle **DELETE** requests appropriately. For more information, see [Restricting Access to Amazon S3 Content by Using an Origin Access Identity \(p. 209\)](#).
- **If you're using a custom origin:** Configure your origin server to handle all methods. For example, if you configure CloudFront to accept and forward these methods *only* because you want to use **POST**, you must still configure your origin server to handle **DELETE** requests appropriately.

Field Level Encryption Config

If you want to enforce field-level encryption on specific data fields, in the drop-down list, choose a field-level encryption configuration.

For more information, see [Using Field-Level Encryption to Help Protect Sensitive Data \(p. 219\)](#).

Cached HTTP Methods

Specify whether you want CloudFront to cache the response from your origin when a viewer submits an `OPTIONS` request. CloudFront always caches the response to `GET` and `HEAD` requests.

Cache Based on Selected Request Headers

Specify whether you want CloudFront to cache objects based on the values of specified headers:

- **None (improves caching)** – CloudFront doesn't cache your objects based on header values.
- **Whitelist** – CloudFront caches your objects based only on the values of the specified headers. Use **Whitelist Headers** to choose the headers that you want CloudFront to base caching on.
- **All** – CloudFront doesn't cache the objects that are associated with this cache behavior. Instead, CloudFront sends every request to the origin. (Not recommended for Amazon S3 origins.)

Regardless of the option that you choose, CloudFront forwards certain headers to your origin and takes specific actions based on the headers that you forward. For more information about how CloudFront handles header forwarding, see [HTTP Request Headers and CloudFront Behavior \(Custom and S3 Origins\)](#) (p. 273).

For more information about how to configure caching in CloudFront by using request headers, see [Caching Content Based on Request Headers](#) (p. 246).

Whitelist Headers

Specify the headers that you want CloudFront to consider when caching your objects. Select headers from the list of available headers and choose **Add**. To forward a custom header, enter the name of the header in the field, and choose **Add Custom**.

For the current maximum number of headers that you can whitelist for each cache behavior, or to request a higher quota (formerly known as limit), see [Quotas on Custom Headers \(Web Distributions Only\)](#) (p. 499).

Object Caching

If your origin server is adding a `Cache-Control` header to your objects to control how long the objects stay in the CloudFront cache and if you don't want to change the `Cache-Control` value, choose **Use Origin Cache Headers**.

To specify a minimum and maximum time that your objects stay in the CloudFront cache regardless of `Cache-Control` headers, and a default time that your objects stay in the CloudFront cache when the `Cache-Control` header is missing from an object, choose **Customize**. Then specify values in the **Minimum TTL**, **Default TTL**, and **Maximum TTL** fields.

For more information, see [Managing How Long Content Stays in an Edge Cache \(Expiration\)](#) (p. 236).

Minimum TTL

Specify the minimum amount of time, in seconds, that you want objects to stay in CloudFront caches before CloudFront forwards another request to your origin to determine whether the object has been updated. The default value for **Minimum TTL** is 0 seconds.

Important

If you configure CloudFront to forward all headers to your origin for a cache behavior, CloudFront never caches the associated objects. Instead, CloudFront forwards all requests for those objects to the origin. In that configuration, the value of **Minimum TTL** must be 0.

To specify a value for **Minimum TTL**, you must choose the **Customize** option for the **Object Caching** setting.

For more information, see [Managing How Long Content Stays in an Edge Cache \(Expiration\)](#) (p. 236).

Maximum TTL

Specify the maximum amount of time, in seconds, that you want objects to stay in CloudFront caches before CloudFront queries your origin to see whether the object has been updated. The value that you specify for **Maximum TTL** applies only when your origin adds HTTP headers such as `Cache-Control max-age`, `Cache-Control s-maxage`, or `Expires` to objects. For more information, see [Managing How Long Content Stays in an Edge Cache \(Expiration\)](#) (p. 236).

To specify a value for **Maximum TTL**, you must choose the **Customize** option for the **Object Caching** setting.

The default value for **Maximum TTL** is 31536000 seconds (one year). If you change the value of **Minimum TTL** or **Default TTL** to more than 31536000 seconds, then the default value of **Maximum TTL** changes to the value of **Default TTL**.

Default TTL

Specify the default amount of time, in seconds, that you want objects to stay in CloudFront caches before CloudFront forwards another request to your origin to determine whether the object has been updated. The value that you specify for **Default TTL** applies only when your origin does *not* add HTTP headers such as `Cache-Control max-age`, `Cache-Control s-maxage`, or `Expires` to objects. For more information, see [Managing How Long Content Stays in an Edge Cache \(Expiration\)](#) (p. 236).

To specify a value for **Default TTL**, you must choose the **Customize** option for the **Object Caching** setting.

The default value for **Default TTL** is 86400 seconds (one day). If you change the value of **Minimum TTL** to more than 86400 seconds, then the default value of **Default TTL** changes to the value of **Minimum TTL**.

Forward Cookies

Note

This option applies to only Amazon S3 buckets that are configured as a website endpoint.

Specify whether you want CloudFront to forward cookies to your origin server and, if so, which ones. If you choose to forward only selected cookies (a whitelist of cookies), enter the cookie names in the **Whitelist Cookies** field. If you choose **All**, CloudFront forwards all cookies regardless of how many your application uses.

Amazon S3 doesn't process cookies, and forwarding cookies to the origin reduces cacheability. For cache behaviors that are forwarding requests to an Amazon S3 origin, choose **None** for **Forward Cookies**.

For more information about forwarding cookies to the origin, go to [Caching Content Based on Cookies](#) (p. 244).

Whitelist Cookies

Note

This option applies to only Amazon S3 buckets that are configured as a website endpoint.

If you chose **Whitelist** in the **Forward Cookies** list, then in the **Whitelist Cookies** field, enter the names of cookies that you want CloudFront to forward to your origin server for this cache behavior. Enter each cookie name on a new line.

You can specify the following wildcards to specify cookie names:

- `*` matches 0 or more characters in the cookie name
- `?` matches exactly one character in the cookie name

For example, suppose viewer requests for an object include a cookie named:

`userid_`*member-number*

Where each of your users has a unique value for *member-number*. You want CloudFront to cache a separate version of the object for each member. You could accomplish this by forwarding all cookies to your origin, but viewer requests include some cookies that you don't want CloudFront to cache. Alternatively, you could specify the following value as a cookie name, which causes CloudFront to forward to the origin all of the cookies that begin with `userid_`:

`userid_*`

For the current maximum number of cookie names that you can whitelist for each cache behavior, or to request a higher quota (formerly known as limit), see [Quotas on Whitelisted Cookies \(Web Distributions Only\)](#) (p. 498).

Query String Forwarding and Caching

CloudFront can cache different versions of your content based on the values of query string parameters. Choose one of the following options:

None (Improves Caching)

Choose this option if your origin returns the same version of an object regardless of the values of query string parameters. This increases the likelihood that CloudFront can serve a request from the cache, which improves performance and reduces the load on your origin.

Forward all, cache based on whitelist

Choose this option if your origin server returns different versions of your objects based on one or more query string parameters. Then specify the parameters that you want CloudFront to use as a basis for caching in the [Query String Whitelist](#) (p. 52) field.

Forward all, cache based on all

Choose this option if your origin server returns different versions of your objects for all query string parameters.

For more information about caching based on query string parameters, including how to improve performance, see [Caching Content Based on Query String Parameters](#) (p. 241).

Query String Whitelist

If you chose **Forward all, cache based on whitelist** for [Query String Forwarding and Caching](#) (p. 52), specify the query string parameters that you want CloudFront to use as a basis for caching.

Smooth Streaming

Choose **Yes** if you want to distribute media files in the Microsoft Smooth Streaming format and you do not have an IIS server.

Choose **No** if you have a Microsoft IIS server that you want to use as an origin to distribute media files in the Microsoft Smooth Streaming format, or if you are not distributing Smooth Streaming media files.

Note

If you specify **Yes**, you can still distribute other content using this cache behavior if that content matches the value of **Path Pattern**.

For more information, see [Configuring Video on Demand for Microsoft Smooth Streaming \(p. 298\)](#).

Restrict Viewer Access (Use Signed URLs or Signed Cookies)

If you want requests for objects that match the `PathPattern` for this cache behavior to use public URLs, choose **No**.

If you want requests for objects that match the `PathPattern` for this cache behavior to use signed URLs, choose **Yes**. Then specify the AWS accounts that you want to use to create signed URLs; these accounts are known as trusted signers.

For more information about trusted signers, see [Specifying the AWS Accounts That Can Create Signed URLs and Signed Cookies \(Trusted Signers\) \(p. 150\)](#).

Trusted Signers

Choose which AWS accounts you want to use as trusted signers for this cache behavior:

- **Self:** Use the account with which you're currently signed into the AWS Management Console as a trusted signer. If you're currently signed in as an IAM user, the associated AWS account is added as a trusted signer.
- **Specify Accounts:** Enter account numbers for trusted signers in the **AWS Account Numbers** field.

To create signed URLs, an AWS account must have at least one active CloudFront key pair.

Important

If you're updating a distribution that you're already using to distribute content, add trusted signers only when you're ready to start generating signed URLs for your objects. After you add trusted signers to a distribution, users must use signed URLs to access the objects that match the `PathPattern` for this cache behavior.

AWS Account Numbers

If you want to create signed URLs using AWS accounts in addition to or instead of the current account, enter one AWS account number per line in this field. Note the following:

- The accounts that you specify must have at least one active CloudFront key pair. For more information, see [Creating CloudFront Key Pairs for Your Trusted Signers \(p. 151\)](#).
- You can't create CloudFront key pairs for IAM users, so you can't use IAM users as trusted signers.
- For information about how to get the AWS account number for an account, see [How Do I Get Security Credentials?](#) in the *Amazon Web Services General Reference*.
- If you enter the account number for the current account, CloudFront automatically checks the **Self** check box and removes the account number from the **AWS Account Numbers** list.

Compress Objects Automatically

If you want CloudFront to automatically compress files of certain types when viewers support compressed content, choose **Yes**. When CloudFront compresses your content, downloads are faster because the files are smaller, and your webpages render faster for your users. For more information, see [Serving compressed files \(p. 116\)](#).

CloudFront Event

You can choose to run a Lambda function when one or more of the following CloudFront events occur:

- When CloudFront receives a request from a viewer (viewer request)
- Before CloudFront forwards a request to the origin (origin request)
- When CloudFront receives a response from the origin (origin response)
- Before CloudFront returns the response to the viewer (viewer response)

For more information, see [How to Decide Which CloudFront Event to Use to Trigger a Lambda Function \(p. 341\)](#).

Lambda Function ARN

Specify the Amazon Resource Name (ARN) of the Lambda function that you want to add a trigger for. To learn how to get the ARN for a function, see step 1 of the procedure [Adding Triggers by Using the CloudFront Console](#).

Distribution Settings

The following values apply to the entire distribution.

Price Class

Choose the price class that corresponds with the maximum price that you want to pay for CloudFront service. By default, CloudFront serves your objects from edge locations in all CloudFront Regions.

For more information about price classes and about how your choice of price class affects CloudFront performance for your distribution, see [Choosing the price class for a CloudFront distribution \(p. 10\)](#). For information about CloudFront pricing, including how price classes map to CloudFront Regions, go to [Amazon CloudFront Pricing](#).

AWS WAF Web ACL

If you want to use AWS WAF to allow or block requests based on criteria that you specify, choose the web ACL to associate with this distribution.

AWS WAF is a web application firewall that lets you monitor the HTTP and HTTPS requests that are forwarded to CloudFront, and lets you control access to your content. Based on conditions that you specify, such as the IP addresses that requests originate from or the values of query strings, CloudFront responds to requests either with the requested content or with an HTTP 403 status code (Forbidden). You can also configure CloudFront to return a custom error page when a request is blocked. For more information about AWS WAF, see the [AWS WAF Developer Guide](#).

Alternate Domain Names (CNAMEs)

Optional. Specify one or more domain names that you want to use for URLs for your objects instead of the domain name that CloudFront assigns when you create your distribution. You must own the domain name, or have authorization to use it, which you verify by adding an SSL/TLS certificate.

For example, if you want the URL for the object:

```
/images/image.jpg
```

To look like this:

`http://www.example.com/images/image.jpg`

Instead of like this:

`http://d111111abcdef8.cloudfront.net/images/image.jpg`

Add a CNAME for `www.example.com`.

Important

If you add a CNAME for `www.example.com` to your distribution, you also must do the following:

- Create (or update) a CNAME record with your DNS service to route queries for `www.example.com` to `d111111abcdef8.cloudfront.net`.
- Add a certificate to CloudFront from a trusted certificate authority (CA) that covers the domain name (CNAME) that you add to your distribution, to validate your authorization to use the domain name.

You must have permission to create a CNAME record with the DNS service provider for the domain. Typically, this means that you own the domain, or that you're developing an application for the domain owner.

For the current maximum number of alternate domain names that you can add to a distribution, or to request a higher quota (formerly known as limit), see [General Quotas on Web Distributions \(p. 497\)](#).

For more information about alternate domain names, see [Using Custom URLs for Files by Adding Alternate Domain Names \(CNAMEs\) \(p. 71\)](#). For more information about CloudFront URLs, see [Customizing the URL Format for Files in CloudFront \(p. 106\)](#).

SSL Certificate

If you specified an alternate domain name to use with your distribution, choose **Custom SSL Certificate**, and then, to validate your authorization to use the alternate domain name, choose a certificate that covers it. If you want viewers to use HTTPS to access your objects, choose the settings that support that.

Note

Before you can specify a custom SSL certificate, you must specify a valid alternate domain name. For more information, see [Requirements for Using Alternate Domain Names \(p. 78\)](#) and [Using Alternate Domain Names and HTTPS \(p. 132\)](#).

- **Default CloudFront Certificate (*.cloudfront.net)** – Choose this option if you want to use the CloudFront domain name in the URLs for your objects, such as `https://d111111abcdef8.cloudfront.net/image1.jpg`.
- **Custom SSL Certificate** – Choose this option if you want to use your own domain name in the URLs for your objects as an alternate domain name, such as `https://example.com/image1.jpg`. Then choose a certificate to use that covers the alternate domain name. The list of certificates can include any of the following:
 - Certificates provided by AWS Certificate Manager
 - Certificates that you purchased from a third-party certificate authority and uploaded to ACM
 - Certificates that you purchased from a third-party certificate authority and uploaded to the IAM certificate store

If you choose this setting, we recommend that you use only an alternate domain name in your object URLs (`https://example.com/logo.jpg`). If you use your CloudFront distribution domain name (`https://d111111abcdef8.cloudfront.net/logo.jpg`) and a client uses an older viewer that doesn't support SNI, how the viewer responds depends on the value that you choose for **Clients Supported**:

- **All Clients**: The viewer displays a warning because the CloudFront domain name doesn't match the domain name in your SSL/TLS certificate.

- **Only Clients that Support Server Name Indication (SNI):** CloudFront drops the connection with the viewer without returning the object.

Custom SSL Client Support

If you specified one or more alternate domain names and a custom SSL certificate for the distribution, choose how you want CloudFront to serve HTTPS requests:

- **Clients that Support Server Name Indication (SNI) - (Recommended)** – With this setting, virtually all modern web browsers and clients can connect to the distribution, because they support SNI. However, some viewers might use older web browsers or clients that don't support SNI, which means they can't connect to the distribution.

To apply this setting using the CloudFront API, specify `sni-only` in the `SSLSupportMethod` field. In AWS CloudFormation, the field is named `SslSupportMethod` (note the different capitalization).

- **Legacy Clients Support** – With this setting, older web browsers and clients that don't support SNI can connect to the distribution. However, this setting incurs additional monthly charges. For the exact price, go to the [Amazon CloudFront Pricing](#) page, and search the page for **Dedicated IP custom SSL**.

To apply this setting using the CloudFront API, specify `vip` in the `SSLSupportMethod` field. In AWS CloudFormation, the field is named `SslSupportMethod` (note the different capitalization).

For more information, see [Choosing How CloudFront Serves HTTPS Requests \(p. 132\)](#).

Security Policy

Specify the security policy that you want CloudFront to use for HTTPS connections with viewers (clients). A security policy determines two settings:

- The minimum SSL/TLS protocol that CloudFront uses to communicate with viewers.
- The ciphers that CloudFront can use to encrypt the content that it returns to viewers.

For more information about the security policies, including the protocols and ciphers that each one includes, see [Supported protocols and ciphers between viewers and CloudFront \(p. 128\)](#).

The security policies that are available depend on the values that you specify for **SSL Certificate** and **Custom SSL Client Support** (known as `CloudFrontDefaultCertificate` and `SSLSupportMethod` in the CloudFront API):

- When **SSL Certificate** is **Default CloudFront Certificate (*.cloudfront.net)** (when `CloudFrontDefaultCertificate` is `true` in the API), CloudFront automatically sets the security policy to `TLSv1`.
- When **SSL Certificate** is **Custom SSL Certificate (example.com)** and **Custom SSL Client Support** is **Clients that Support Server Name Indication (SNI) - (Recommended)** (when `CloudFrontDefaultCertificate` is `false` and `SSLSupportMethod` is `sni-only` in the API), you can choose from the following security policies:
 - `TLSv1.2_2019`
 - `TLSv1.2_2018`
 - `TLSv1.1_2016`
 - `TLSv1_2016`
 - `TLSv1`
- When **SSL Certificate** is **Custom SSL Certificate (example.com)** and **Custom SSL Client Support** is **Legacy Clients Support** (when `CloudFrontDefaultCertificate` is `false` and `SSLSupportMethod` is `vip` in the API), you can choose from the following security policies:

- TLSv1
- SSLv3

In this configuration, the TLSv1.2_2019, TLSv1.2_2018, TLSv1.1_2016, and TLSv1_2016 security policies aren't available in the CloudFront console or API. If you want to use one of these security policies, you have the following options:

- Evaluate whether your distribution needs Legacy Clients Support with dedicated IP addresses. If your viewers support [server name indication \(SNI\)](#), we recommend that you update your distribution's **Custom SSL Client Support** setting to **Clients that Support Server Name Indication (SNI)** (set `SSLSupportMethod` to `sni-only` in the API). This enables you to use any of the available TLS security policies, and it can also reduce your CloudFront charges.
- If you must keep Legacy Clients Support with dedicated IP addresses, you can request one of the other TLS security policies (TLSv1.2_2019, TLSv1.2_2018, TLSv1.1_2016, or TLSv1_2016) by creating a case in the [AWS Support Center](#).

Note

Before you contact AWS Support to request this change, consider the following:

- When you add one of these security policies (TLSv1.2_2019, TLSv1.2_2018, TLSv1.1_2016, or TLSv1_2016) to a Legacy Clients Support distribution, the security policy is applied to *all* non-SNI viewer requests for *all* Legacy Clients Support distributions in your AWS account. However, when viewers send SNI requests to a distribution with Legacy Clients Support, the security policy of that distribution applies. To make sure that your desired security policy is applied to *all* viewer requests sent to *all* Legacy Clients Support distributions in your AWS account, add the desired security policy to each distribution individually.
- By definition, the new security policy doesn't support the same ciphers and protocols as the old one. For example, if you chose to upgrade a distribution's security policy from TLSv1 to TLSv1.1_2016, that distribution will no longer support the DES-CBC3-SHA cipher. For more information about the ciphers and protocols that each security policy supports, see [Supported protocols and ciphers between viewers and CloudFront \(p. 128\)](#).

Supported HTTP Versions

Choose the HTTP versions that you want your distribution to support when viewers communicate with CloudFront. For viewers and CloudFront to use HTTP/2, viewers must support TLS 1.2 or later, and must support Server Name Identification (SNI).

In general, configuring CloudFront to communicate with viewers using HTTP/2 reduces latency. You can improve performance by optimizing for HTTP/2. For more information, do an internet search for "http/2 optimization."

Default Root Object

Optional. The object that you want CloudFront to request from your origin (for example, `index.html`) when a viewer requests the root URL of your distribution (`http://www.example.com/`) instead of an object in your distribution (`http://www.example.com/product-description.html`). Specifying a default root object avoids exposing the contents of your distribution.

The maximum length of the name is 255 characters. The name can contain any of the following characters:

- A-Z, a-z
- 0-9
- `_ - . * $ / ~ ' "`

- `&`, passed and returned as `&`;

When you specify the default root object, enter only the object name, for example, `index.html`. Do not add a `/` before the object name.

For more information, see [Specifying a Default Root Object](#) (p. 107).

Logging

Whether you want CloudFront to log information about each request for an object and store the log files in an Amazon S3 bucket. You can enable or disable logging at any time. There is no extra charge if you enable logging, but you accrue the usual Amazon S3 charges for storing and accessing the files in an Amazon S3 bucket. You can delete the logs at any time. For more information about CloudFront access logs, see [Configuring and Using Standard Logs \(Access Logs\)](#) (p. 439).

Bucket for Logs

If you chose **On** for **Logging**, the Amazon S3 bucket that you want CloudFront to store access logs in, for example, `myLogs-DOC-EXAMPLE-BUCKET.s3.amazonaws.com`.

Note

Don't choose an Amazon S3 bucket in any of the following Regions, because CloudFront doesn't deliver access logs to buckets in these Regions:

- Africa (Cape Town) `af-south-1`
- Asia Pacific (Hong Kong) `ap-east-1`
- Europe (Milan) `eu-south-1`
- Middle East (Bahrain) `me-south-1`

The [Amazon S3 console](#) shows the bucket's Region.

If you enable logging, CloudFront records information about each end-user request for an object and stores the files in the specified Amazon S3 bucket. You can enable or disable logging at any time. For more information about CloudFront access logs, see [Configuring and Using Standard Logs \(Access Logs\)](#) (p. 439).

Note

You must have the permissions required to get and update Amazon S3 bucket ACLs, and the S3 ACL for the bucket must grant you `FULL_CONTROL`. This allows CloudFront to give the `awsdatafeeds` account permission to save log files in the bucket. For more information, see [Permissions Required to Configure Standard Logging and to Access Your Log Files](#) (p. 440).

Log Prefix

Optional. If you chose **On** for **Logging**, specify the string, if any, that you want CloudFront to prefix to the access log file names for this distribution, for example, `exampleprefix/`. The trailing slash (`/`) is optional but recommended to simplify browsing your log files. For more information about CloudFront access logs, see [Configuring and Using Standard Logs \(Access Logs\)](#) (p. 439).

Cookie Logging

If you want CloudFront to include cookies in access logs, choose **On**. If you choose to include cookies in logs, CloudFront logs all cookies regardless of how you configure the cache behaviors for this distribution: forward all cookies, forward no cookies, or forward a specified list of cookies to the origin.

Amazon S3 doesn't process cookies, so unless your distribution also includes an Amazon EC2 or other custom origin, we recommend that you choose **Off** for the value of **Cookie Logging**.

For more information about cookies, go to [Caching Content Based on Cookies \(p. 244\)](#).

Enable IPv6

IPv6 is a new version of the IP protocol. It's the eventual replacement for IPv4 and uses a larger address space. CloudFront always responds to IPv4 requests. If you want CloudFront to respond to requests from IPv4 IP addresses (such as 192.0.2.44) and requests from IPv6 addresses (such as 2001:0db8:85a3:0000:0000:8a2e:0370:7334), select **Enable IPv6**.

In general, you should enable IPv6 if you have users on IPv6 networks who want to access your content. However, if you're using signed URLs or signed cookies to restrict access to your content, and if you're using a custom policy that includes the `IpAddress` parameter to restrict the IP addresses that can access your content, do not enable IPv6. If you want to restrict access to some content by IP address and not restrict access to other content (or restrict access but not by IP address), you can create two distributions. For information about creating signed URLs by using a custom policy, see [Creating a Signed URL Using a Custom Policy \(p. 166\)](#). For information about creating signed cookies by using a custom policy, see [Setting Signed Cookies Using a Custom Policy \(p. 182\)](#).

If you're using a Route 53 alias resource record set to route traffic to your CloudFront distribution, you need to create a second alias resource record set when both of the following are true:

- You enable IPv6 for the distribution
- You're using alternate domain names in the URLs for your objects

For more information, see [Routing Traffic to an Amazon CloudFront Web Distribution by Using Your Domain Name](#) in the *Amazon Route 53 Developer Guide*.

If you created a CNAME resource record set, either with Route 53 or with another DNS service, you don't need to make any changes. A CNAME record routes traffic to your distribution regardless of the IP address format of the viewer request.

If you enable IPv6 and CloudFront access logs, the `c-ip` column includes values in IPv4 and IPv6 format. For more information, see [Configuring and Using Standard Logs \(Access Logs\) \(p. 439\)](#).

Note

To maintain high customer availability, CloudFront responds to viewer requests by using IPv4 if our data suggests that IPv4 will provide a better user experience. To find out what percentage of requests CloudFront is serving over IPv6, enable CloudFront logging for your distribution and parse the `c-ip` column, which contains the IP address of the viewer that made the request. This percentage should grow over time, but it will remain a minority of traffic as IPv6 is not yet supported by all viewer networks globally. Some viewer networks have excellent IPv6 support, but others don't support IPv6 at all. (A viewer network is analogous to your home internet or wireless carrier.)

For more information about our support for IPv6, see the [CloudFront FAQ](#). For information about enabling access logs, see the fields [Logging \(p. 58\)](#), [Bucket for Logs \(p. 58\)](#), and [Log Prefix \(p. 58\)](#).

Comment

Optional. When you create a distribution, you can include a comment of up to 128 characters. You can update the comment at any time.

Distribution State

Indicates whether you want the distribution to be enabled or disabled once it's deployed:

- *Enabled* means that as soon as the distribution is fully deployed you can deploy links that use the distribution's domain name and users can retrieve content. Whenever a distribution is enabled,

CloudFront accepts and handles any end-user requests for content that use the domain name associated with that distribution.

When you create, modify, or delete a CloudFront distribution, it takes time for your changes to propagate to the CloudFront database. An immediate request for information about a distribution might not show the change. Propagation usually completes within minutes, but a high system load or network partition might increase this time.

- *Disabled* means that even though the distribution might be deployed and ready to use, users can't use it. Whenever a distribution is disabled, CloudFront doesn't accept any end-user requests that use the domain name associated with that distribution. Until you switch the distribution from disabled to enabled (by updating the distribution's configuration), no one can use it.

You can toggle a distribution between disabled and enabled as often as you want. Follow the process for updating a distribution's configuration. For more information, see [Updating a Distribution \(p. 63\)](#).

Custom Error Pages and Error Caching

You can have CloudFront return an object to the viewer (for example, an HTML file) when your Amazon S3 or custom origin returns an HTTP 4xx or 5xx status code to CloudFront. You can also specify how long an error response from your origin or a custom error page is cached in CloudFront edge caches. For more information, see [Creating a Custom Error Page for Specific HTTP Status Codes \(p. 291\)](#).

Note

The following values aren't included in the Create Distribution wizard, so you can configure custom error pages only when you update a distribution.

HTTP Error Code

The HTTP status code for which you want CloudFront to return a custom error page. You can configure CloudFront to return custom error pages for none, some, or all of the HTTP status codes that CloudFront caches.

Error Caching Minimum TTL (seconds)

The minimum amount of time that you want CloudFront to cache error responses from your origin server.

Response Page Path

The path to the custom error page (for example, `/4xx-errors/403-forbidden.html`) that you want CloudFront to return to a viewer when your origin returns the HTTP status code that you specified for **Error Code** (for example, 403). If you want to store your objects and your custom error pages in different locations, your distribution must include a cache behavior for which the following is true:

- The value of **Path Pattern** matches the path to your custom error messages. For example, suppose you saved custom error pages for 4xx errors in an Amazon S3 bucket in a directory named `/4xx-errors`. Your distribution must include a cache behavior for which the path pattern routes requests for your custom error pages to that location, for example, `/4xx-errors/*`.
- The value of **Origin** specifies the value of **Origin ID** for the origin that contains your custom error pages.

HTTP Response Code

The HTTP status code that you want CloudFront to return to the viewer along with the custom error page.

Restrictions

If you need to prevent users in selected countries from accessing your content, you can configure your CloudFront distribution either to allow users in a whitelist of specified countries to access your content or to not allow users in a blacklist of specified countries to access your content. For more information, see [Restricting the Geographic Distribution of Your Content \(p. 217\)](#).

Note

The following values aren't included in the Create Distribution wizard, so you can configure geo restrictions only when you update a distribution.

Enable Geo-Restriction

Whether you want to prevent users in selected countries from accessing your content. There is no additional charge for configuring geo restriction.

Restriction Type

How you want to specify the countries from which your users can access your content:

- **Whitelist:** The **Countries** list includes all of the countries from which you *do* want your users to access your content.
- **Blacklist:** The **Countries** list includes all of the countries from which you *do not* want your users to access your content.

Countries

The countries that you want to add to your whitelist or blacklist. To add a country, select it in the list on the left and choose **Add**. Note the following:

- To add multiple consecutive countries, select the first country, press and hold the Shift key, select the last country, and choose **Add**.
- To add multiple non-consecutive countries, select the first country, press and hold the Ctrl key, select the remaining countries, and choose **Add**.
- To find a country in the left list, enter the first few characters of the country's full name.
- The two-letter code before the name of each country is the value that you enter if you want to create or update a distribution by using the CloudFront API. We use the International Organization for Standardization country codes. For an easy-to-use list, sortable by code and by country name, see the Wikipedia entry [ISO 3166-1 alpha-2](#).

Values That CloudFront Displays in the Console

When you create a new distribution or update an existing distribution, CloudFront displays the following information in the CloudFront console.

Note

Active trusted signers, the AWS accounts that have an active CloudFront key pair and can be used to create valid signed URLs, are currently not visible in the CloudFront console.

Distribution ID (General Tab)

When you perform an action on a distribution using the CloudFront API, you use the distribution ID to specify which distribution to use, for example, `EDFDVBBD6EXAMPLE`. You can't change a distribution's distribution ID.

Distribution Status (General Tab)

The possible status values for a distribution are listed in the following table.

Value	Description
InProgress	The distribution is still being created or updated, and the changes have not yet fully propagated to edge servers.
Deployed	The distribution has been created or updated and the changes have been fully propagated through the CloudFront system.

Note

In addition to ensuring that the status for a distribution is **Deployed**, you must enable the distribution before users can use CloudFront to access your content. For more information, see [Distribution State \(p. 59\)](#).

Last Modified (General Tab)

The date and time that the distribution was last modified, using ISO 8601 format, for example, 2012-05-19T19:37:58Z. For more information, see <http://www.w3.org/TR/NOTE-datetime>.

Domain Name (General Tab)

You use the distribution's domain name in the links to your objects. For example, if your distribution's domain name is `d111111abcdef8.cloudfront.net`, the link to `/images/image.jpg` would be `http://d111111abcdef8.cloudfront.net/images/image.jpg`. You can't change the CloudFront domain name for your distribution. For more information about CloudFront URLs for links to your objects, see [Customizing the URL Format for Files in CloudFront \(p. 106\)](#).

If you specified one or more alternate domain names (CNAMEs), you can use your own domain names for links to your objects instead of using the CloudFront domain name. For more information about CNAMEs, see [Alternate Domain Names \(CNAMEs\) \(p. 54\)](#).

Note

CloudFront domain names are unique. Your distribution's domain name was never used for a previous distribution and will never be reused for another distribution in the future.

Testing a Distribution

After you've created your distribution, CloudFront knows where your origin server is, and you know the domain name associated with the distribution. You can create links to your objects using the CloudFront domain name, and CloudFront will serve the objects to your webpage or application.

Note

You must wait until the status of the distribution changes to **Deployed** before you can test your links.

To create links to objects in a web distribution

1. Copy the following HTML code into a new file, replace `domain-name` with your distribution's domain name, and replace `object-name` with the name of your object.

```
<html>
<head>My CloudFront Test</head>
<body>
<p>My text content goes here.</p>
<p>
```

```
</body>  
</html>
```

For example, if your domain name were `d111111abcdef8.cloudfront.net` and your object were `image.jpg`, the URL for the link would be:

```
http://d111111abcdef8.cloudfront.net/image.jpg.
```

If your object is in a folder on your origin server, then the folder must also be included in the URL. For example, if `image.jpg` were located in the `images` folder on your origin server, then the URL would be:

```
http://d111111abcdef8.cloudfront.net/images/image.jpg
```

2. Save the HTML code in a file that has an `.html` file name extension.
3. Open your webpage in a browser to ensure that you can see your object.

The browser returns your page with the embedded image file, served from the edge location that CloudFront determined was appropriate to serve the object.

Updating a Distribution

In the CloudFront console, you can see the CloudFront distributions that are associated with your AWS account, view the settings for a distribution, and update most settings. Be aware that settings changes that you make won't take effect until the distribution has propagated to the AWS edge locations.

To Update a CloudFront Distribution

1. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
2. Select the ID of a distribution. The list includes all of the distributions associated with the AWS account that you used to sign in to the CloudFront console.
3. To edit settings for a distribution, choose the **Distribution Settings** tab.
4. To make updates, do one of the following:
 - a. To update general settings, choose **Edit**. Otherwise, choose the tab for the settings that you want to update: **Origins** or **Behaviors**.
 - b. For settings for an RTMP distribution, choose **Edit**, and then update values.

For information about the fields, see [Values that You Specify When You Create or Update an RTMP Distribution \(p. 309\)](#).
5. Make the updates, and then, to save your changes, choose **Yes, Edit**. For information about the fields, see the following topics:
 - **General settings:** [Distribution Settings \(p. 54\)](#)
 - **Origin settings:** [Origin Settings \(p. 40\)](#)
 - **Cache behavior settings:** [Cache Behavior Settings \(p. 46\)](#)
6. If you want to delete an origin in your distribution, do the following:
 - a. Choose **Behaviors**, and then make sure you have moved any default cache behaviors associated with the origin to another origin.
 - b. Choose **Origins**, and then select an origin.
 - c. Choose **Delete**.

You can also update a distribution by using the CloudFront API:

- To update a distribution, see [UpdateDistribution](#) in the *Amazon CloudFront API Reference*.

Important

When you update your distribution, be aware that a number of additional fields are required that are not required to create a distribution. For a summary of the fields required for when you create or update a distribution, see [Required fields for creating and updating distributions \(p. 34\)](#). To help make sure that all of the required fields are included when you update a distribution by using the CloudFront API, follow the steps described in [UpdateDistribution](#) in the *Amazon CloudFront API Reference*.

When you save changes to your distribution configuration, CloudFront starts to propagate the changes to all edge locations. Until your configuration is updated in an edge location, CloudFront continues to serve your content from that location based on the previous configuration. After your configuration is updated in an edge location, CloudFront immediately starts to serve your content from that location based on the new configuration.

Your changes don't propagate to every edge location instantaneously. When propagation is complete, the status of your distribution changes from **InProgress** to **Deployed**. While CloudFront is propagating your changes, we unfortunately can't determine whether a given edge location is serving your content based on the previous configuration or the new configuration.

Tagging Amazon CloudFront Distributions

Tags are words or phrases that you can use to identify and organize your AWS resources. You can add multiple tags to each resource, and each tag includes a key and a value that you define. For example, the key might be "domain" and the value might be "example.com". You can search and filter your resources based on the tags you add.

The following are two examples of how it can be useful to work with tags in CloudFront:

- Use tags to track billing information in different categories. When you apply tags to CloudFront distributions or other AWS resources (such as Amazon EC2 instances or Amazon S3 buckets) and activate the tags, AWS generates a cost allocation report as a comma-separated value (CSV file) with your usage and costs aggregated by your active tags. You can apply tags that represent business categories (such as cost centers, application names, or owners) to organize your costs across multiple services. For more information about using tags for cost allocation, see [Use Cost Allocation Tags](#) in the *AWS Billing and Cost Management User Guide*.
- Use tags to enforce tag-based permissions on CloudFront distributions. For more information, see [Tag-Based Policies \(p. 480\)](#).

Note the following:

- You can tag web and RTMP distributions, but you can't tag origin access identities or invalidations.
- [Tag Editor](#) and [Resource Groups](#) are currently not supported for CloudFront.

For the current maximum number of tags that you can add to a distribution, see [Quotas \(p. 496\)](#). To request a higher quota (formerly known as limit), [create a case](#) with the AWS Support Center.

You can also apply tags to resources by using the CloudFront API, AWS CLI, SDKs, and AWS Tools for Windows PowerShell. For more information, see the following documentation:

- CloudFront API – See the following operations in the *Amazon CloudFront API Reference*:
 - [ListTagsForResource](#)
 - [TagResource](#)
 - [UntagResource](#)

- AWS CLI – See [cloudfront](#) in the *AWS CLI Command Reference*
- SDKs – See the applicable SDK documentation on the [AWS Documentation](#) page
- Tools for Windows PowerShell – See [Amazon CloudFront](#) in the *AWS Tools for PowerShell Cmdlet Reference*

Topics

- [Tag Restrictions \(p. 65\)](#)
- [Adding, Editing, and Deleting Tags for Distributions \(p. 65\)](#)

Tag Restrictions

The following basic restrictions apply to tags:

- Maximum number of tags per resource – 50
- Maximum key length – 128 Unicode characters
- Maximum value length – 256 Unicode characters
- Valid values for key and value – a-z, A-Z, 0-9, space, and the following characters: _ . : / = + - and @
- Tag keys and values are case sensitive
- Don't use `aws :` as a prefix for keys; it's reserved for AWS use

Adding, Editing, and Deleting Tags for Distributions

The following procedure explains how to add, edit, and delete tags for your distributions in the CloudFront console.

To add tags, edit, or delete tags for a distribution

1. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
2. Choose the ID for the distribution that you want to update.
3. Choose the **Tags** tab.
4. Choose **Add or edit tags**.
5. On the Add or edit tags page, you can do the following:

Add a tag

Enter a key and, optionally, a value for the tag.

Edit a tag

Change the key, the value, or both. You can also delete the value for a tag, but the key is required.

Delete a tag

Choose the **X** on the right side of the value field.

6. Choose **Save**.

Deleting a Distribution

If you no longer want to use a distribution, you can delete it by using the CloudFront console or by using the CloudFront API.

Be aware that before you can delete a distribution, you must disable it, which requires permission to update the distribution. For more information about setting permissions for working with CloudFront, including setting `UpdateDistribution` and `DeleteDistribution` permissions, see [Customer Managed Policy Examples](#) (p. 485).

Note

If you disable a distribution that has an alternate domain name associated with it, CloudFront stops accepting traffic for that domain name (such as `www.example.com`), even if another distribution has an alternate domain name with a wildcard (*) that matches the same domain (such as `*.example.com`).

To Delete a CloudFront Distribution

1. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
2. In the right pane of the CloudFront console, find the distribution that you want to delete.
3. If the value of the **State** column is **Disabled**, skip to Step 7.

If the value of **State** is **Enabled** and the value of **Status** is **Deployed**, continue with Step 4 to disable the distribution before deleting it.

If the value of **State** is **Enabled** and the value of **Status** is **InProgress**, wait until **Status** changes to **Deployed**. Then continue with Step 4 to disable the distribution before deleting it.

4. In the right pane of the CloudFront console, select the check box for the distribution that you want to delete.
5. Choose **Disable** to disable the distribution, and choose **Yes, Disable** to confirm. Then choose **Close**.

Note

Because CloudFront must propagate this change to all edge locations, it might take a few minutes before the update is complete and you can delete your distribution.

6. The value of the **State** column immediately changes to **Disabled**. Wait until the value of the **Status** column changes to **Deployed**.
7. Check the check box for the distribution that you want to delete.
8. Choose **Delete**, and choose **Yes, Delete** to confirm. Then click **Close**.

Note

If you have just marked your distribution as disabled, CloudFront might still need a few more minutes to propagate that change to the edge locations. Until propagation is complete, the **Delete** option isn't available.

You can also delete a distribution using the CloudFront API:

- To delete a distribution, see [DeleteDistribution](#) in the *Amazon CloudFront API Reference*.
- To delete an RTMP distribution, see [DeleteStreamingDistribution](#) in the *Amazon CloudFront API Reference*.

Using Amazon S3 Origins, MediaPackage Channels, and Custom Origins for Web Distributions

When you create a distribution, you specify where CloudFront sends requests for the files. CloudFront supports using several AWS resources as origins. For example, you can specify an Amazon S3 bucket or a MediaStore container, a MediaPackage channel, or a custom origin, such as an Amazon EC2 instance or your own HTTP web server.

Topics

- [Using Amazon S3 Buckets for Your Origin \(p. 67\)](#)
- [Using Amazon S3 Buckets Configured as Website Endpoints for Your Origin \(p. 68\)](#)
- [Using a MediaStore Container or a MediaPackage Channel for Your Origin \(p. 68\)](#)
- [Using Amazon EC2 or Other Custom Origins \(p. 68\)](#)
- [Using CloudFront Origin Groups \(p. 69\)](#)
- [Adding CloudFront When You're Already Distributing Content from Amazon S3 \(p. 69\)](#)
- [Moving an Amazon S3 Bucket to a Different Region \(p. 71\)](#)

Using Amazon S3 Buckets for Your Origin

When you use Amazon S3 as an origin for your distribution, you place any objects that you want CloudFront to deliver in an Amazon S3 bucket. You can use any method that is supported by Amazon S3 to get your objects into Amazon S3, for example, the Amazon S3 console or API, or a third-party tool. You can create a hierarchy in your bucket to store the objects, just as you would with any other Amazon S3 bucket.

Using an existing Amazon S3 bucket as your CloudFront origin server doesn't change the bucket in any way; you can still use it as you normally would to store and access Amazon S3 objects at the standard Amazon S3 price. You incur regular Amazon S3 charges for storing the objects in the bucket. For more information about the charges to use CloudFront, see [CloudFront Reports in the Console \(p. 405\)](#).

Important

For your bucket to work with CloudFront, the name must conform to DNS naming requirements. For more information, go to [Bucket Restrictions and Limitations](#) in the *Amazon Simple Storage Service Developer Guide*.

When you specify the Amazon S3 bucket that you want CloudFront to get objects from, we recommend that you use the following format to access the bucket:

`bucket-name.s3.region.amazonaws.com`

Another option might be to use the following more general format, but be aware that this format doesn't work for Regions launched in 2019 or later:

`bucket-name.s3.amazonaws.com`

When you specify the bucket name in this format, you can use the following CloudFront features:

- Configure CloudFront to communicate with your Amazon S3 bucket using SSL. For more information, see [Using HTTPS with CloudFront \(p. 121\)](#).
- Use an origin access identity to require that your users access your content using CloudFront URLs, not by using Amazon S3 URLs. For more information, see [Restricting Access to Amazon S3 Content by Using an Origin Access Identity \(p. 209\)](#).
- Update the content of your bucket by submitting `POST` and `PUT` requests to CloudFront. For more information, see [HTTP Methods \(p. 265\)](#) in the topic [How CloudFront Processes and Forwards Requests to Your Amazon S3 Origin Server \(p. 263\)](#).

Do not specify the bucket using the following formats:

- The Amazon S3 path style, `s3.amazonaws.com/bucket-name`
- The Amazon S3 CNAME, if any

Using Amazon S3 Buckets Configured as Website Endpoints for Your Origin

You can set up an Amazon S3 bucket that is configured as a website endpoint as custom origin with CloudFront.

- When you configure your CloudFront distribution, for the origin, enter the Amazon S3 static website hosting endpoint for your bucket. This value appears in the Amazon S3 console, on the **Properties** tab, in the **Static website hosting** pane. For example:

`http://bucket-name.s3-website-region.amazonaws.com`

For more information about specifying Amazon S3 static website endpoints, see [Website endpoints](#) in the *Amazon Simple Storage Service Developer Guide*.

When you specify the bucket name in this format as your origin, you can use Amazon S3 redirects and Amazon S3 custom error documents. For more information about Amazon S3 features, see the [Amazon S3 documentation](#). (CloudFront also provides custom error pages. For more information, see [Creating a Custom Error Page for Specific HTTP Status Codes \(p. 291\)](#).)

Using an Amazon S3 bucket as your CloudFront origin server doesn't change it in any way. You can still use it as you normally would and you incur regular Amazon S3 charges. For more information about the charges to use CloudFront, see [CloudFront Pricing](#).

Note

If you use the CloudFront API to create your distribution with an Amazon S3 bucket that is configured as a website endpoint, you must configure it by using `CustomOriginConfig`, even though the website is hosted in an Amazon S3 bucket. For more information about creating distributions by using the CloudFront API, see [CreateDistribution](#) in the *Amazon CloudFront API Reference*.

Using a MediaStore Container or a MediaPackage Channel for Your Origin

To stream video by using CloudFront, you can set up an Amazon S3 bucket that is configured as a MediaStore container, or create a channel and endpoints with MediaPackage. Then you create and configure a distribution in CloudFront to stream the video.

For more information and step-by-step instructions, see the following topics:

- [Serving Video Using AWS Elemental MediaStore as the Origin \(p. 300\)](#)
- [Serving Live Video Formatted with AWS Elemental MediaPackage \(p. 301\)](#)

Using Amazon EC2 or Other Custom Origins

A custom origin is an HTTP server, for example, a web server. The HTTP server can be an Amazon Elastic Compute Cloud (Amazon EC2) instance or an HTTP server that you manage privately. An Amazon S3 origin configured as a website endpoint is also considered a custom origin.

When you use a custom origin that is your own HTTP server, you specify the DNS name of the server, along with the HTTP and HTTPS ports and the protocol that you want CloudFront to use when fetching objects from your origin.

Most CloudFront features are supported when you use a custom origin with the following exceptions:

- **RTMP distributions**—Not supported.
- **Private content**—Although you can use a signed URL to distribute content from a custom origin, for CloudFront to access the custom origin, the origin must remain publicly accessible. For more information, see [Serving Private Content with Signed URLs and Signed Cookies \(p. 145\)](#).

Follow these guidelines for using Amazon EC2 instances and other custom origins with CloudFront.

- Host and serve the same content on all servers that are serving content for the same CloudFront origin. For more information, see [Origin Settings \(p. 40\)](#) in the [Values That You Specify When You Create or Update a Distribution \(p. 38\)](#) topic.
- Log the `X-Amz-Cf-Id` header entries on all servers; CloudFront requires this information for debugging.
- Restrict access requests to the HTTP and HTTPS ports that your custom origin listens on.
- Synchronize the clocks of all servers in your implementation. Note that CloudFront uses Coordinated Universal Time (UTC) for signed URLs and signed cookies, for access logs, and reports. In addition, if you monitor CloudFront activity using CloudWatch metrics, note that CloudWatch also uses UTC.
- Use redundant servers to handle failures.
- For information about using a custom origin to serve private content, see [Restricting Access to Files on Custom Origins \(p. 148\)](#).
- For information about request and response behavior and about supported HTTP status codes, see [Request and Response Behavior \(p. 263\)](#).

If you use Amazon EC2 for your custom origins, we recommend that you do the following:

1. Use an Amazon Machine Image that automatically installs the software for a web server. For more information, see the [Amazon EC2 documentation](#).
2. Use an Elastic Load Balancing load balancer to handle traffic across multiple Amazon EC2 instances and to isolate your application from changes to Amazon EC2 instances. For example, if you use a load balancer, you can add and delete Amazon EC2 instances without changing your application. For more information, see the [Elastic Load Balancing documentation](#).
3. When you create your CloudFront distribution, specify the URL of the load balancer for the domain name of your origin server. For more information, see [Creating a Distribution \(p. 37\)](#).

Using CloudFront Origin Groups

You can specify an origin group for your CloudFront origin if, for example, you want to configure origin failover for scenarios when you need high availability. Use origin failover to designate a primary origin for CloudFront plus a second origin that CloudFront automatically switches to when the primary origin returns specific HTTP status code failure responses.

To see the steps for setting up an origin group and for more information, see [Optimizing High Availability with CloudFront Origin Failover \(p. 231\)](#).

Adding CloudFront When You're Already Distributing Content from Amazon S3

If you store your objects in an Amazon S3 bucket, you can either have users get your objects directly from S3, or you can configure CloudFront to get your objects from S3 and then distribute them to your users.

Note

To learn more about using Amazon S3 buckets for your origin with CloudFront, including when you have an Amazon S3 bucket configured as a website endpoint, see [Using Amazon S3 Origins, MediaPackage Channels, and Custom Origins for Web Distributions](#) (p. 66).

Using CloudFront can be more cost effective if your users access your objects frequently because, at higher usage, the price for CloudFront data transfer is lower than the price for Amazon S3 data transfer. In addition, downloads are faster with CloudFront than with Amazon S3 alone because your objects are stored closer to your users.

Note

If you want CloudFront to respect Amazon S3 cross-origin resource sharing settings, configure CloudFront to forward the `Origin` header to Amazon S3. For more information, see [Caching Content Based on Request Headers](#) (p. 246).

If you currently distribute content directly from your Amazon S3 bucket using your own domain name (such as `example.com`) instead of the domain name of your Amazon S3 bucket (such as `MyAWSBucket.s3.us-west-2.amazonaws.com`), you can add CloudFront with no disruption by using the following procedure.

To add CloudFront when you're already distributing your content from Amazon S3

1. Create a CloudFront distribution using one of the following procedures:

- [Steps for Creating a Distribution \(Overview\)](#) (p. 36)
- [Task List for Streaming Media Files Using RTMP](#) (p. 307)

When you create the distribution, specify the name of your Amazon S3 bucket as the origin server.

Important

For your bucket to work with CloudFront, the name must conform to DNS naming requirements. For more information, see [Bucket Restrictions and Limitations](#) in the *Amazon Simple Storage Service Developer Guide*.

If you're using a CNAME with Amazon S3, specify the CNAME for your distribution, too.

2. Create a test webpage that contains links to publicly readable objects in your Amazon S3 bucket, and test the links. For this initial test, use the CloudFront domain name of your distribution in the object URLs, for example, `http://d111111abcdef8.cloudfront.net/images/image.jpg`.

For more information about the format of CloudFront URLs, see [Customizing the URL Format for Files in CloudFront](#) (p. 106).

3. If you're using Amazon S3 CNAMEs, your application uses your domain name (for example, `example.com`) to reference the objects in your Amazon S3 bucket instead of using the name of your bucket (for example, `DOC-EXAMPLE-BUCKET1.s3.amazonaws.com`). To continue using your domain name to reference objects instead of using the CloudFront domain name for your distribution (for example, `d111111abcdef8.cloudfront.net`), you need to update your settings with your DNS service provider.

For Amazon S3 CNAMEs to work, your DNS service provider must have a CNAME resource record set for your domain that currently routes queries for the domain to your Amazon S3 bucket. For example, if a user requests this object:

```
http://example.com/images/image.jpg
```

The request is automatically rerouted, and the user sees this object:

```
http://DOC-EXAMPLE-BUCKET1.s3.amazonaws.com/images/image.jpg
```

To route queries to your CloudFront distribution instead of your Amazon S3 bucket, you need to use the method provided by your DNS service provider to update the CNAME resource record set for your domain. This updated CNAME record starts to redirect DNS queries from your domain to the CloudFront domain name for your distribution. For more information, see the documentation provided by your DNS service provider.

Note

If you're using Route 53 as your DNS service, you can use either a CNAME resource record set or an alias resource record set. For information about editing resource record sets, see [Editing Resource Record Sets](#). For information about alias resource record sets, see [Choosing Between Alias and Non-Alias Resource Record Sets](#). Both topics are in the *Amazon Route 53 Developer Guide*.

For more information about using CNAMEs with CloudFront, see [Using Custom URLs for Files by Adding Alternate Domain Names \(CNAMEs\)](#) (p. 71).

After you update the CNAME resource record set, it can take up to 72 hours for the change to propagate throughout the DNS system, although it usually happens faster. During this time, some requests for your content will continue to be routed to your Amazon S3 bucket, and others will be routed to CloudFront.

Moving an Amazon S3 Bucket to a Different Region

If you're using Amazon S3 as the origin for a CloudFront distribution and you move the bucket to a different Region, CloudFront can take up to an hour to update its records to include the change of Region when both of the following are true:

- You're using a CloudFront origin access identity (OAI) to restrict access to the bucket.
- You move the bucket to an Amazon S3 Region that requires Signature Version 4 for authentication.

When you're using OAIs, CloudFront uses the Region (among other values) to calculate the signature that it uses to request objects from your bucket. For more information about OAIs, see [Restricting Access to Amazon S3 Content by Using an Origin Access Identity](#) (p. 209). For a list of Amazon S3 Regions and the signature versions that they support, see [Amazon Simple Storage Service \(Amazon S3\)](#) in the "Regions and Endpoints" chapter of the *Amazon Web Services General Reference*.

To force a faster update to CloudFront's records, you can update your CloudFront distribution, for example, by updating the **Comment** field on the **General** tab in the CloudFront console. When you update a distribution, CloudFront immediately checks on the Region that your bucket is in; propagation of the change to all edge locations should take less than 15 minutes.

Using Custom URLs for Files by Adding Alternate Domain Names (CNAMEs)

In CloudFront, an alternate domain name, also known as a CNAME, lets you use your own domain name (for example, `www.example.com`) in your files' URLs instead of using the domain name that CloudFront assigns to your distribution. Both web and RTMP distributions support alternate domain names.

When you create a distribution, CloudFront returns a domain name for the distribution, for example:

```
d1111111abcdef8.cloudfront.net
```


When you use the CloudFront domain name for your files, the URL for a file called `/images/image.jpg` is:

```
https://d1111111abcdef8.cloudfront.net/images/image.jpg
```

If you want to use your own domain name, such as `www.example.com`, instead of the `cloudfront.net` domain name, you can add an alternate domain name to your distribution, like `www.example.com`. You can then use the following URL to view `/images/image.jpg`:

```
https://www.example.com/images/image.jpg
```

Topics

- [Adding an Alternate Domain Name \(p. 72\)](#)
- [Moving an Alternate Domain Name to a Different CloudFront Distribution \(p. 74\)](#)
- [Removing an Alternate Domain Name \(p. 77\)](#)
- [Using Wildcards in Alternate Domain Names That You Add to CloudFront \(p. 78\)](#)
- [Requirements for Using Alternate Domain Names \(p. 78\)](#)
- [Restrictions on Using Alternate Domain Names \(p. 80\)](#)

Adding an Alternate Domain Name

The following task list describes how to use the CloudFront console to add an alternate domain name to your distribution so that you can use your own domain name in your links instead of the CloudFront domain name. For information about updating your distribution using the CloudFront API, see [Working with distributions \(p. 33\)](#).

Note

If you want viewers to use HTTPS with your alternate domain name, see [Using Alternate Domain Names and HTTPS \(p. 132\)](#).

Before you begin: Make sure that you do the following before you update your distribution to add an alternate domain name:

- Register the domain name with Route 53 or another domain provider.
- Add a certificate from an authorized certificate authority (CA) to CloudFront that covers the domain name you plan to use with the distribution, to validate that you are authorized to use the domain. For more information, see [Requirements for Using Alternate Domain Names \(p. 78\)](#).

Adding an Alternate Domain Name

1. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
2. Choose the ID for the distribution that you want to update.
3. On the **General** tab, choose **Edit**.
4. Update the following values:

Alternate Domain Names (CNAMEs)

Add your alternate domain names. Separate domain names with commas, or type each domain name on a new line.

SSL Certificate (Web Distributions Only)

Choose the following setting:

- **Use HTTPS** – Choose **Custom SSL Certificate**, and then choose a certificate from the list. The list includes certificates provisioned by AWS Certificate Manager (ACM), certificates that you purchased from another CA and uploaded to ACM, and certificates that you purchased from another CA and uploaded to the IAM certificate store.

If you uploaded a certificate to the IAM certificate store but it doesn't appear in the list, review the procedure [Importing an SSL/TLS Certificate \(p. 140\)](#) to confirm that you correctly uploaded the certificate.

If you choose this setting, we recommend that you use only an alternate domain name in your object URLs (<https://www.example.com/logo.jpg>). If you use your CloudFront distribution domain name (<https://d111111abcdef8.cloudfront.net/logo.jpg>), a viewer might behave as follows, depending on the value that you choose for **Clients Supported**:

- **All Clients**: If the viewer doesn't support SNI, it displays a warning because the CloudFront domain name doesn't match the domain name in your TLS/SSL certificate.
- **Only Clients that Support Server Name Indication (SNI)**: CloudFront drops the connection with the viewer without returning the object.

Clients Supported (Web Distributions Only)

Choose an option:

- **All Clients**: CloudFront serves your HTTPS content using dedicated IP addresses. If you select this option, you incur additional charges when you associate your SSL/TLS certificate with a distribution that is enabled. For more information, see [Amazon CloudFront Pricing](#).
- **Only Clients that Support Server Name Indication (SNI) (Recommended)**: Older browsers or other clients that don't support SNI must use another method to access your content.

For more information, see [Choosing How CloudFront Serves HTTPS Requests \(p. 132\)](#).

5. Choose **Yes, Edit**.
6. On the **General** tab for the distribution, confirm that **Distribution Status** has changed to **Deployed**. If you try to use an alternate domain name before the updates to your distribution have been deployed, the links that you create in the following steps might not work.
7. Configure the DNS service for the domain to route traffic for the domain, such as www.example.com, to the CloudFront domain name for your distribution, such as d111111abcdef8.cloudfront.net. The method that you use depends on whether you're using Route 53 as the DNS service provider for the domain or another provider.

Note

If your DNS record already points to a distribution that is not the distribution that you are updating, then you only add the alternate domain name to your distribution after you update your DNS. For more information, see [Restrictions on Using Alternate Domain Names \(p. 80\)](#).

Route 53

Create an alias resource record set. With an alias resource record set, you don't pay for Route 53 queries. In addition, you can create an alias resource record set for the root domain name (example.com), which DNS doesn't allow for CNAMEs. For more information, see [Routing traffic to an Amazon CloudFront web distribution by using your domain name](#) in the *Amazon Route 53 Developer Guide*.

Another DNS service provider

Use the method provided by your DNS service provider to add a CNAME record for your domain. This new CNAME record will redirect DNS queries from your domain (for example, www.example.com) to the CloudFront domain name for your distribution (for example,

d111111abcdef8.cloudfront.net). For more information, see the documentation provided by your DNS service provider.

Important

If you already have an existing CNAME record for your domain name, update that record or replace it with a new one that points to the CloudFront domain name for your distribution.

In addition, confirm that your CNAME resource record set points to your distribution's domain name and not to one of your origin servers.

8. Using `dig` or a similar DNS tool, confirm that the resource record set that you created in step 7 points to the domain name for your distribution.

The following example shows a `dig` request on the `www.example.com` domain, as well as the relevant part of the response.

```
PROMPT> dig www.example.com

; <<> DiG 9.3.3rc2 <<> www.example.com
;; global options: printcmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 15917
;; flags: qr rd ra; QUERY: 1, ANSWER: 9, AUTHORITY: 2, ADDITIONAL: 0

;; QUESTION SECTION:
;www.example.com.      IN      A

;; ANSWER SECTION:
www.example.com. 10800 IN CNAME d111111abcdef8.cloudfront.net.
...
```

The answer section shows a CNAME record that routes queries for `www.example.com` to the CloudFront distribution domain name `d111111abcdef8.cloudfront.net`. If the name on the right side of CNAME is the domain name for your CloudFront distribution, the CNAME record is configured correctly. If that is any other value, for example, the domain name for your Amazon S3 bucket, then the CNAME record is configured incorrectly. In that case, go back to step 4 and correct the CNAME record to point to the domain name for your distribution.

9. Test the alternate domain name by visiting URLs with your domain name instead of the CloudFront domain name for your distribution.
10. In your application, change the URLs for your objects to use your alternate domain name instead of the domain name of your CloudFront distribution.

Moving an Alternate Domain Name to a Different CloudFront Distribution

If you want to move an alternate domain name from one CloudFront distribution to another distribution, the steps you must take depend on the domain name that you want to move:

- For a subdomain like `www.example.com`, you can move the domain yourself. For detailed steps, see [Move a subdomain name to another distribution \(p. 74\)](#).
- For a domain like `example.com` (an apex domain), you must work with AWS Support to move the domain to another distribution [Move an apex domain name to another distribution \(p. 76\)](#).

Move a subdomain name to another distribution

Follow these steps to move a subdomain name, for example `www.example.com`.

To move a subdomain name to a new distribution

1. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
2. If you don't have a new distribution to move the domain name to, create one. For more information, see [Creating a Distribution](#) (p. 37).
3. Add to the distribution an alternate domain name that includes a wildcard for the alias record set or CNAME record. For example, if the subdomain name that you want to move to the new distribution is `www.example.com`, add the alternate domain name `*.example.com`. For more information, see [Using Wildcards in Alternate Domain Names That You Add to CloudFront](#) (p. 78).

Note

You can't add a wildcard to a top-level domain name, such as `*.com`, so if you want to move a domain name like `example.com` to a new distribution, see [Move an apex domain name to another distribution](#) (p. 76).

4. Update the DNS configuration for your subdomain to point to the new distribution. For example, you would update the DNS service for the subdomain `www.example.com` to route traffic to the CloudFront domain name for your distribution, `d111111abcdef8.cloudfront.net`.

To update the configuration, do one of the following:

- **If you're using Route 53**, update alias records or CNAME records, depending how you set up the alternate domain name originally. For more information, see [Editing records](#) in the *Amazon Route 53 Developer Guide*.
- **If you're using another DNS service provider**, use the method provided by the DNS service provider to update the CNAME record that directs traffic to CloudFront. For more information, see the documentation provided by your DNS service provider.

Note

At this point, the subdomain is still served by the original distribution because that is where the alternate domain is currently configured.

5. Using `dig` or a similar DNS tool, confirm that the resource record set that you created in step 4 points to the domain name for your distribution.

The following example shows a `dig` request on the `www.example.com` domain, as well as the relevant part of the response.

```
PROMPT> dig www.example.com

; <<> DiG 9.3.3rc2 <<> www.example.com
;; global options: printcmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 15917
;; flags: qr rd ra; QUERY: 1, ANSWER: 9, AUTHORITY: 2, ADDITIONAL: 0

;; QUESTION SECTION:
;www.example.com.      IN      A

;; ANSWER SECTION:
www.example.com. 10800 IN CNAME d111111abcdef8.cloudfront.net.
...
```

The answer section shows a CNAME record that routes queries for `www.example.com` to the CloudFront distribution domain name `d111111abcdef8.cloudfront.net`. If the name on the right side of CNAME is the domain name for your CloudFront distribution, the CNAME record is configured correctly. If that is any other value, for example, the domain name for your Amazon S3 bucket, then the CNAME record is configured incorrectly. In that case, go back to step 4 and correct the CNAME record to point to the domain name for your distribution.

Note

When using Route 53 alias records, it is not possible to use `dig` to confirm that the resource record points to the new distribution. In this case, you can either change the type of the resource record from alias to CNAME, or wait until the record's time to live (TTL) value has expired.

6. Remove the CNAME from the existing distribution and then add it to the new CloudFront distribution where the wildcard alternate name was added previously.

Note

While these changes propagate, the alternate domain name may be served by either the original or new distribution at random. This behavior may persist for a few minutes after both distributions have reached status **Deployed**.

7. Test the alternate domain name by visiting some URLs that use your domain name instead of the CloudFront domain name for your distribution.
8. (Optional) The wildcard alternate name may now be removed from the new distribution.
9. If you're no longer using the original distribution, delete it. For more information, see [Deleting a Distribution](#) (p. 65).

Move an apex domain name to another distribution

For apex domain names, like `example.com`, you must contact AWS Support to move the domain name to another CloudFront distribution. The extra steps are required because moving a domain yourself, as described in the previous procedure, requires setting up domain routing using a wildcard for part of the domain name. For apex domains, for this step, you would have to set up routing as `*.com`, which isn't allowed.

Before you get started, if you don't have a new distribution to move the domain name to, create one. For more information, see [Creating a Distribution](#) (p. 37).

Moving a domain name like `example.com` to a new distribution takes two steps:

Step 1: Provide proof to AWS Support that you own the domain name by creating a TXT record for your domain at your DNS service provider. This helps prevent someone else from making changes to your distribution configuration.

Step 2: Request that AWS Support move your domain to the new CloudFront distribution.

Here are the specific steps to take.

Step 1: Create a TXT record for your domain

1. Sign in to your DNS service provider website.

If your service provider is Route 53, [sign in to the Route 53 console](#).

2. Create a TXT record for your domain like the following:

`domain name` TXT `CloudFrontdistribution domain name`

For example: `example.com` TXT `d1111111abcdef8.cloudfront.net`

- If your DNS service provider is Route 53, go to step 3 for detailed steps.
- If your domain is hosted by another DNS service provider, see the documentation at the DNS service provider. You might need to request that your service provider create the TXT record for you.

Tip

If your service provider does not allow a TXT record for a domain to have the same information as a CNAME record, consider creating a TXT record that uses your domain

name with an underscore (_) prepended to it. For an example, see the following Knowledge Center article: [Resolve CNAME Already Exists Error](#).

3. If your DNS service provider is Route 53, use the following steps to create a TXT record to prove domain ownership:
 - a. On the **Hosted Zones** page, double-click the row for the hosted zone in which you want to edit records.
 - b. Choose **Create Record Set**.
 - c. Enter the following values:
 - **Name:** The domain name you want to move to a new CloudFront distribution
 - **Type:** TXT
 - **Alias:** No
 - **TTL:** 60 seconds
 - **Value:** The name of the CloudFront distribution that you want to add this domain name to, such as d123.cloudfront.net
 - **Routing policy:** Simple
 - d. Choose **Create**.

Step 2: Request that AWS Support move your domain to the new CloudFront distribution

- Sign in to AWS and [contact AWS support](#) to request that they verify that you own the domain, and move the domain to the new CloudFront distribution.

Note

AWS Support can't verify your domain ownership until they can view the TXT record that you created for your domain. Be aware that records that you create at your DNS provider can take a while (up to several days) to propagate through the DNS system.

Removing an Alternate Domain Name

If you want to stop routing traffic for a domain or subdomain to a CloudFront distribution, follow the steps in this section to update both the DNS configuration and the CloudFront distribution.

It's important that you remove the alternate domain names from the distribution as well as update your DNS configuration. This helps prevent issues later if you want to associate the domain name with another CloudFront distribution. If an alternate domain name is already associated with one distribution, it can't be set up with another.

Note

If you want to remove the alternate domain name from this distribution so you can add it to another one, follow the steps in [Moving an Alternate Domain Name to a Different CloudFront Distribution \(p. 74\)](#). If you follow the steps here instead (to remove a domain) and then add the domain to another distribution, there will be a period of time during which the domain won't link to the new distribution because CloudFront is propagating the updates to edge locations.

To remove an alternate domain name from a distribution

1. To start, route internet traffic for your domain to another resource that isn't your CloudFront distribution, such as an Elastic Load Balancing load balancer. Or you can delete the DNS record that's routing traffic to CloudFront.

Do one of the following, depending on the DNS service for your domain:

- **If you're using Route 53**, update or delete alias records or CNAME records. For more information, see [Editing records](#) or [Deleting records](#).
 - **If you're using another DNS service provider**, use the method provided by the DNS service provider to update or delete the CNAME record that directs traffic to CloudFront. For more information, see the documentation provided by your DNS service provider.
2. After you update your domain's DNS records, wait until the changes have propagated and DNS resolvers are routing traffic to the new resource. You can check to see when this is complete by creating some test links that use your domain in the URL.
 3. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>, and update your CloudFront distribution to remove the domain name by doing the following:
 - a. Choose the ID for the distribution that you want to update.
 - b. On the **General** tab, choose **Edit**.
 - c. In **Alternate Domain Names (CNAMEs)**, remove the alternate domain name (or domain names) that you no longer want to use for your distribution.
 - d. Choose **Yes, Edit**.

Using Wildcards in Alternate Domain Names That You Add to CloudFront

When you add alternate domain names, you can use the * wildcard at the beginning of a domain name instead of adding subdomains individually. For example, with an alternate domain name of *.example.com, you can use any domain name that ends with example.com in your object URLs, such as www.example.com, product-name.example.com, and marketing.product-name.example.com. The name of an object is the same regardless of the domain name, for example:

```
www.example.com/images/image.jpg
```

```
product-name.example.com/images/image.jpg
```

```
marketing.product-name.example.com/images/image.jpg
```

Follow these requirements for alternate domain names that include wildcards:

- The alternate domain name must begin with an asterisk and a dot (* .).
- You *cannot* use a wildcard to replace part of a subdomain name, like this: *domain.example.com.
- You cannot replace a subdomain in the middle of a domain name, like this: subdomain.*.example.com.
- All alternate domain names, including alternate domain names that use wildcards, must be covered by the subject alternative name (SAN) on the certificate.

A wildcard alternate domain name, such as *.example.com, can include another alternate domain name, such as example.com, as long as they're both in the same CloudFront distribution or they're in distributions that were created by using the same AWS account.

Requirements for Using Alternate Domain Names

When you add an alternate domain name, such as www.example.com, to a CloudFront distribution, the following are requirements:

Alternate domain names must be lowercase

All alternate domain names (CNAMEs) must be lowercase to be valid.

Alternate domain names must be covered by a valid SSL/TLS certificate

To add an alternate domain name (CNAME) to use with a CloudFront distribution, you must attach to your distribution a trusted, valid SSL/TLS certificate that covers the alternate domain name. This ensures that only people with access to your domain's certificate can associate with CloudFront a CNAME related to your domain.

A trusted certificate is one that is issued by AWS Certificate Manager (ACM) or by another valid certificate authority (CA); you can't use a self-signed certificate. CloudFront supports the same certificate authorities as Mozilla. For the current list, see [Mozilla Included CA Certificate List](#).

To verify an alternate domain name by using the certificate that you attach, including alternate domain names that include wildcards, CloudFront checks the subject alternative name (SAN) on the certificate. The alternate domain name that you're adding must be covered by the SAN.

Note

Only one certificate can be attached to a CloudFront distribution at a time.

You prove that you are authorized to add a specific alternate domain name to your distribution by doing one of the following:

- Attaching a certificate that includes the alternate domain name, like `product-name.example.com`.
- Attaching a certificate that includes a `*` wildcard at the beginning of a domain name, to cover multiple subdomains with one certificate. When you specify a wildcard, you can add multiple subdomains as alternate domain names in CloudFront.

The following examples illustrate how using wildcards in domain names in a certificate work to authorize you to add specific alternate domain names in CloudFront.

- You want to add `marketing.example.com` as an alternate domain name. You list in your certificate the following domain name: `*.example.com`. When you attach this certificate to CloudFront, you can add any alternate domain name for your distribution that replaces the wildcard at that level, including `marketing.example.com`. You can also, for example, add the following alternate domain names:
 - `product.example.com`
 - `api.example.com`

However, you can't add alternate domain names that are at levels higher or lower than the wildcard. For example, you can't add the alternate domain names `example.com` or `marketing.product.example.com`.

- You want to add `example.com` as an alternate domain name. To do this, you must list the domain name `example.com` itself on the certificate that you attach to your distribution.
- You want to add `marketing.product.example.com` as an alternate domain name. To do this, you can list `*.product.example.com` on the certificate, or you can list `marketing.product.example.com` itself on the certificate.

Permission to change DNS configuration

When you add alternate domain names, you must create CNAME records to route DNS queries for the domain names to your CloudFront distribution. To do this, you must have permission to create CNAME records with the DNS service provider for the alternate domain names that you're using. Typically, this means that you own the domains, but you might be developing an application for the domain owner.

Alternate Domain Names and HTTPS

If you want viewers to use HTTPS with an alternate domain name, you must complete some additional configuration. For more information, see [Using Alternate Domain Names and HTTPS \(p. 132\)](#).

Restrictions on Using Alternate Domain Names

Note the following restrictions on using alternate domain names:

Maximum number of alternate domain names

For the current maximum number of alternate domain names that you can add to a distribution, or to request a higher quota (formerly known as limit), see [General Quotas on Web Distributions \(p. 497\)](#).

Duplicate and overlapping alternate domain names

You cannot add an alternate domain name to a CloudFront distribution if the alternate domain name already exists in another CloudFront distribution, even if your AWS account owns the other distribution.

However, you can add a wildcard alternate domain name, such as `*.example.com`, that includes (that overlaps with) a non-wildcard alternate domain name, such as `www.example.com`. Overlapping alternate domain names can be in the same distribution or in separate distributions as long as both distributions were created by using the same AWS account.

If you have overlapping alternate domain names in two distributions, CloudFront sends the request to the distribution with the more specific name match, regardless of the distribution that the DNS record points to. For example, `marketing.domain.com` is more specific than `*.domain.com`.

Alternate domain names that already point to a distribution

If your DNS record points to a distribution that is not the distribution that you are creating or modifying, then you can't add the alternate domain name to your distribution. In this scenario, you must update your DNS at your DNS provider before you can add the domain name for your CloudFront distribution.

To correct this, sign in to your DNS provider and remove the existing DNS record, or contact your DNS provider to remove it for you. Then create the correct DNS record for your distribution, following the steps for adding or changing the alternate domain name for a distribution. For more information, see [Adding an Alternate Domain Name \(p. 72\)](#) or [Moving an Alternate Domain Name to a Different CloudFront Distribution \(p. 74\)](#).

Domain fronting

CloudFront includes protection against domain fronting occurring across different AWS accounts. Domain fronting is a scenario in which a non-standard client creates a TLS/SSL connection to a domain name in one AWS account, but then makes an HTTPS request for an unrelated name in another AWS account. For example, the TLS connection might connect to `www.example.com`, and then issue a request for `www.example.org`.

To prevent cases where domain fronting crosses different AWS accounts, CloudFront makes sure that the AWS account that owns the certificate that it serves for a specific connection always matches the AWS account that owns the request that it handles on that same connection.

If the two AWS account numbers do not match, CloudFront responds with an HTTP 421 Misdirected Request response to give the client a chance to connect using the correct domain.

Adding an alternate domain name at the top node (zone apex) for a domain

When you add an alternate domain name to a distribution, you typically create a CNAME record in your DNS configuration to route DNS queries for the domain name to your CloudFront distribution. However, you can't create a CNAME record for the top node of a DNS namespace, also known as the zone apex; the DNS protocol doesn't allow it. For example, if you register the DNS name `example.com`, the zone apex is `example.com`. You can't create a CNAME record for `example.com`, but you can create CNAME records for `www.example.com`, `newproduct.example.com`, and so on.

If you're using Route 53 as your DNS service, you can create an alias resource record set, which has two advantages over CNAME records. You can create an alias resource record set for a domain name at the top node (`example.com`). In addition, when you use an alias resource record set, you don't pay for Route 53 queries.

Note

If you enable IPv6, you must create two alias resource record sets: one to route IPv4 traffic (an A record) and one to route IPv6 traffic (an AAAA record). For more information, see [Enable IPv6 \(p. 59\)](#) in the topic [Values That You Specify When You Create or Update a Distribution \(p. 38\)](#).

For more information, see [Routing traffic to an Amazon CloudFront web distribution by using your domain name](#) in the *Amazon Route 53 Developer Guide*.

Using WebSocket with CloudFront Distributions

Amazon CloudFront supports using WebSocket, a TCP-based protocol that is useful when you need long-lived bidirectional connections between clients and servers. A persistent connection is often a requirement with real-time applications. The scenarios in which you might use Websockets include social chat platforms, online collaboration workspaces, multi-player gaming, and services that provide real-time data feeds like financial trading platforms. Data over a WebSocket connection can flow in both directions for full-duplex communication.

CloudFront supports WebSocket connections globally with no required additional configuration. All CloudFront distributions have built-in WebSocket protocol support, as long as the client and server also both support the protocol.

How the WebSocket Protocol Works

The WebSocket protocol is an independent, TCP-based protocol that allows you to avoid some of the overhead—and potentially increased latency—of HTTP.

To establish a WebSocket connection, the client sends a regular HTTP request that uses HTTP's upgrade semantics to change the protocol. The server can then complete the handshake. The WebSocket connection remains open and either the client or server can send data frames to each other without having to establish new connections each time.

By default, the WebSocket protocol uses port 80 for regular WebSocket connections and port 443 for WebSocket connections over TLS/SSL. The options that you choose for your CloudFront [Viewer Protocol Policy \(p. 49\)](#) and [Origin Protocol Policy \(p. 45\)](#) apply to WebSocket connections as well as to HTTP traffic.

WebSocket Requirements

WebSocket requests must comply with RFC 6455 (<http://tools.ietf.org/html/rfc6455>) in the following standard formats.

Sample client request:

```
GET /chat HTTP/1.1
Host: server.example.com
Upgrade: websocket
Connection: Upgrade
Sec-WebSocket-Key: dGhlIHNhbXBsZSBub25jZQ==
Origin: http://example.com
Sec-WebSocket-Protocol: chat, superchat
Sec-WebSocket-Version: 13
```

Sample server response:

```
HTTP/1.1 101 Switching Protocols
Upgrade: websocket
Connection: Upgrade
Sec-WebSocket-Accept: s3pPLMBiTxaQ9kYGzzhZRbK+xOo=
Sec-WebSocket-Protocol: chat
```

If the WebSocket connection is disconnected by the client or server, or by a network disruption, client applications are expected to re-initiate the connection with the server.

Working with policies

With CloudFront *policies*, you can control the values that are included in the *cache key* for objects that are cached at CloudFront edge locations. These values can include HTTP request query strings, headers, and cookies. The cache key determines whether a viewer request results in a *cache hit* (the object is served to the viewer from a CloudFront edge location).

When there's a *cache miss* (the requested object is not cached at the edge location), CloudFront sends a request to the origin to retrieve the object. This is called an *origin request*. You can separately control which of these values (query strings, headers, and cookies) are included in the origin request.

You control the cache key with a *cache policy* and the origin request with an *origin request policy*. By controlling the cache key and the origin request separately, you can forward request values to your origin without duplicating cached content when the content doesn't differ based on those values.

For more information, see the following.

Contents

- [Controlling the cache key \(p. 83\)](#)
 - [Understanding cache policies \(p. 84\)](#)
 - [Policy information \(p. 84\)](#)
 - [Time to live \(TTL\) settings \(p. 85\)](#)
 - [Cache key settings \(p. 85\)](#)
 - [Creating cache policies \(p. 88\)](#)
 - [Using the managed cache policies \(p. 91\)](#)
 - [Attaching a managed cache policy \(p. 92\)](#)
 - [Understanding the managed cache policies \(p. 92\)](#)
 - [Understanding the cache key \(p. 93\)](#)
 - [The default cache key \(p. 94\)](#)
 - [Customizing the cache key \(p. 94\)](#)
- [Controlling origin requests \(p. 95\)](#)
 - [Understanding origin request policies \(p. 96\)](#)
 - [Policy information \(p. 96\)](#)
 - [Origin request settings \(p. 96\)](#)
 - [Creating origin request policies \(p. 97\)](#)
 - [Using the managed origin request policies \(p. 100\)](#)
 - [Attaching a managed origin request policy \(p. 101\)](#)
 - [Understanding the managed origin request policies \(p. 101\)](#)
- [Using the CloudFront HTTP headers \(p. 102\)](#)
 - [Headers for determining the viewer's device type \(p. 102\)](#)
 - [Headers for determining the viewer's location \(p. 103\)](#)
 - [Other CloudFront headers \(p. 103\)](#)

Controlling the cache key

With Amazon CloudFront, you can control the *cache key* for objects that are cached at CloudFront edge locations. The cache key is the unique identifier for every object in the cache, and it determines whether

a viewer request results in a *cache hit*. A cache hit occurs when a viewer request generates the same cache key as a prior request, and the object for that cache key is in the edge location's cache and valid. When there's a cache hit, the object is served to the viewer from a CloudFront edge location, which has the following benefits:

- Reduced load on your origin server
- Reduced latency for the viewer

You can get better performance from your website or application when you have a higher *cache hit ratio* (a higher proportion of viewer requests result in a cache hit). One way to improve your cache hit ratio is to include only the minimum necessary values in the cache key. For more information, see [Understanding the cache key \(p. 93\)](#).

To control the cache key, you use a CloudFront *cache policy*. You attach a cache policy to one or more cache behaviors in a CloudFront distribution.

Contents

- [Understanding cache policies \(p. 84\)](#)
 - [Policy information \(p. 84\)](#)
 - [Time to live \(TTL\) settings \(p. 85\)](#)
 - [Cache key settings \(p. 85\)](#)
- [Creating cache policies \(p. 88\)](#)
- [Using the managed cache policies \(p. 91\)](#)
 - [Attaching a managed cache policy \(p. 92\)](#)
 - [Understanding the managed cache policies \(p. 92\)](#)
- [Understanding the cache key \(p. 93\)](#)
 - [The default cache key \(p. 94\)](#)
 - [Customizing the cache key \(p. 94\)](#)

Understanding cache policies

You can use a cache policy to improve your cache hit ratio by controlling the values (URL query strings, HTTP headers, and cookies) that are included in the cache key. CloudFront provides some predefined cache policies, known as *managed policies*, for common use cases. You can use these managed policies, or you can create your own cache policy that's specific to your needs. For more information about the managed policies, see [Using the managed cache policies \(p. 91\)](#).

A cache policy contains the following settings, which are categorized into *policy information*, *time to live (TTL) settings*, and *cache key settings*.

Policy information

Name

A name to identify the cache policy. In the console, you use the name to attach the cache policy to a cache behavior.

Comment

A comment to describe the cache policy. This is optional, but it can help you identify the purpose of the cache policy.

Time to live (TTL) settings

The time to live (TTL) settings work together with the `Cache-Control` and `Expires` HTTP headers (if they're in the origin response) to determine how long objects in the CloudFront cache remain valid.

Minimum TTL

The minimum amount of time, in seconds, that you want objects to stay in the CloudFront cache before CloudFront checks with the origin to see if the object has been updated. For more information, see [Managing How Long Content Stays in an Edge Cache \(Expiration\)](#) (p. 236).

Maximum TTL

The maximum amount of time, in seconds, that objects stay in the CloudFront cache before CloudFront checks with the origin to see if the object has been updated. CloudFront uses this setting only when the origin sends `Cache-Control` or `Expires` headers with the object. For more information, see [Managing How Long Content Stays in an Edge Cache \(Expiration\)](#) (p. 236).

Default TTL

The default amount of time, in seconds, that you want objects to stay in the CloudFront cache before CloudFront checks with the origin to see if the object has been updated. CloudFront uses this setting's value as the object's TTL only when the origin does *not* send `Cache-Control` or `Expires` headers with the object. For more information, see [Managing How Long Content Stays in an Edge Cache \(Expiration\)](#) (p. 236).

Cache key settings

Cache key settings specify the values in viewer requests that CloudFront includes in the cache key. The values can include URL query strings, HTTP headers, and cookies. The values that you include in the cache key are automatically included in requests that CloudFront sends to the origin, known as *origin requests*. For information about controlling origin requests without affecting the cache key, see [Controlling origin requests](#) (p. 95).

Cache key settings include:

- [Query strings](#) (p. 85)
- [Headers](#) (p. 86)
- [Cookies](#) (p. 86)
- [Cache compressed objects \(uses the `Accept-Encoding` header\)](#) (p. 86)

Query strings

The URL query strings in viewer requests that CloudFront includes in the cache key and in origin requests. For query strings, you can choose one of the following settings:

- **None** – The query strings in viewer requests are *not* included in the cache key and are *not* automatically included in origin requests.
- **All** – All query strings in viewer requests are included in the cache key and are also automatically included in origin requests.
- **Whitelist** – You specify which of the query strings in viewer requests are included in the cache key and automatically included in origin requests.
- **All-Except** – You specify which of the query strings in viewer requests are *not* included in the cache key and are *not* automatically included in origin requests. All other query strings, except for the ones you specify, *are* included in the cache key and automatically included in origin requests.

When you use the **Whitelist** or **All-Except** setting, you specify query strings by their name, not their value. For example, consider the following URL path:

```
/content/stories/example-story.html?split-pages=false
```

In this case, you specify the query string as `split-pages`, not as `split-pages=false`. However, CloudFront includes the full query string, including its value, in the cache key and in origin requests.

Headers

The HTTP headers in viewer requests that CloudFront includes in the cache key and in origin requests. For headers, you can choose one of the following settings:

- **None** – The HTTP headers in viewer requests are *not* included in the cache key and are *not* automatically included in origin requests.
- **Whitelist** – You specify which of the HTTP headers in viewer requests are included in the cache key and automatically included in origin requests.

When you use the **Whitelist** setting, you specify HTTP headers by their name, not their value. For example, consider the following HTTP header:

```
Accept-Language: en-US,en;q=0.5
```

In this case, you specify the header as `Accept-Language`, not as `Accept-Language: en-US,en;q=0.5`. However, CloudFront includes the full header, including its value, in the cache key and in origin requests.

You can also include certain headers generated by CloudFront in the cache key. For more information, see [Using the CloudFront HTTP headers \(p. 102\)](#).

Cookies

The cookies in viewer requests that CloudFront includes in the cache key and in origin requests. For cookies, you can choose one of the following settings:

- **None** – The cookies in viewer requests are *not* included in the cache key and are *not* automatically included in origin requests.
- **All** – All cookies in viewer requests are included in the cache key and are automatically included in origin requests.
- **Whitelist** – You specify which of the cookies in viewer requests are included in the cache key and automatically included in origin requests.
- **All-Except** – You specify which of the cookies in viewer requests are *not* included in the cache key and are *not* automatically included in origin requests. All other cookies, except for the ones you specify, *are* included in the cache key and automatically included in origin requests.

When you use the **Whitelist** or **All-Except** setting, you specify cookies by their name, not their value. For example, consider the following Cookie header:

```
Cookie: session_ID=abcd1234
```

In this case, you specify the cookie as `session_ID`, not as `session_ID=abcd1234`. However, CloudFront includes the full cookie, including its value, in the cache key and in origin requests.

Cache compressed objects (uses the `Accept-Encoding` header)

These settings enable CloudFront to request and cache objects that are compressed in the Gzip or Brotli compression formats, when the viewer supports it. Viewers indicate their support for these compression formats with the `Accept-Encoding` HTTP header.

Note

The Chrome and Firefox web browsers support Brotli compression only when the request is sent using HTTPS. These browsers do not support Brotli with HTTP requests.

Enable these settings when any of the following are true:

- Your origin returns Gzip compressed objects when viewers support them (requests contain the `Accept-Encoding` HTTP header with `gzip` as a value). In this case, enable the **Cache Gzip objects** setting (set `EnableAcceptEncodingGzip` to `true` in the CloudFront API, AWS SDKs, AWS CLI, or AWS CloudFormation).
- Your origin returns Brotli compressed objects when viewers support them (requests contain the `Accept-Encoding` HTTP header with `br` as a value). In this case, enable the **Cache Brotli objects** setting (set `EnableAcceptEncodingBrotli` to `true` in the CloudFront API, AWS SDKs, AWS CLI, or AWS CloudFormation).
- The cache behavior that this cache policy is attached to is configured with [CloudFront edge compression \(p. 116\)](#). In this case, you can enable caching for either Gzip or Brotli, or both. When CloudFront edge compression is enabled, enabling caching for both formats can help to reduce your costs for data transfer out to the internet.

Note

If you enable caching for one or both of these compression formats, do not include the `Accept-Encoding` header in an [origin request policy \(p. 95\)](#) that's associated with the same cache behavior. CloudFront always includes this header in origin requests when caching is enabled for either of these formats, so including `Accept-Encoding` in an origin request policy has no effect.

If your origin server does not return Gzip or Brotli compressed objects, or the cache behavior is not configured with CloudFront edge compression, don't enable caching for compressed objects. If you do, it might cause a decrease in your [cache hit ratio \(p. 228\)](#).

The following explains how these settings affect a CloudFront distribution. All of the following scenarios assume that the viewer request includes the `Accept-Encoding` header. When the viewer request does not include the `Accept-Encoding` header, CloudFront doesn't include this header in the cache key and doesn't include it in the corresponding origin request.

When caching compressed objects is enabled for both compression formats

If the viewer supports both Gzip and Brotli—that is, if the `gzip` and `br` values are both in the `Accept-Encoding` header in the viewer request—CloudFront does the following:

- Normalizes the header to `Accept-Encoding: br, gzip` and includes the normalized header in the cache key. The cache key doesn't include other values that were in the `Accept-Encoding` header sent by the viewer.
- If the edge location has a Brotli or Gzip compressed object in the cache that matches the request and is not expired, the edge location returns the object to the viewer.
- If the edge location doesn't have a Brotli or Gzip compressed object in the cache that matches the request and is not expired, CloudFront includes the normalized header (`Accept-Encoding: br, gzip`) in the corresponding origin request. The origin request doesn't include other values that were in the `Accept-Encoding` header sent by the viewer.

If the viewer supports one compression format but not the other—for example, if `gzip` is a value in the `Accept-Encoding` header in the viewer request but `br` is not—CloudFront does the following:

- Normalizes the header to `Accept-Encoding: gzip` and includes the normalized header in the cache key. The cache key doesn't include other values that were in the `Accept-Encoding` header sent by the viewer.
- If the edge location has a Gzip compressed object in the cache that matches the request and is not expired, the edge location returns the object to the viewer.

- If the edge location doesn't have a Gzip compressed object in the cache that matches the request and is not expired, CloudFront includes the normalized header (`Accept-Encoding: gzip`) in the corresponding origin request. The origin request doesn't include other values that were in the `Accept-Encoding` header sent by the viewer.

To understand what CloudFront does if the viewer supports Brotli but not Gzip, replace the two compression formats with each other in the preceding example.

If the viewer does not support Brotli or Gzip—that is, the `Accept-Encoding` header in the viewer request does not contain `br` or `gzip` as values—CloudFront:

- Doesn't include the `Accept-Encoding` header in the cache key.
- Includes `Accept-Encoding: identity` in the corresponding origin request. The origin request doesn't include other values that were in the `Accept-Encoding` header sent by the viewer.

When caching compressed objects is enabled for one compression format, but not the other

If the viewer supports the format for which caching is enabled—for example, if caching compressed objects is enabled for Gzip and the viewer supports Gzip (`gzip` is one of the values in the `Accept-Encoding` header in the viewer request)—CloudFront does the following:

- Normalizes the header to `Accept-Encoding: gzip` and includes the normalized header in the cache key.
- If the edge location has a Gzip compressed object in the cache that matches the request and is not expired, the edge location returns the object to the viewer.
- If the edge location doesn't have a Gzip compressed object in the cache that matches the request and is not expired, CloudFront includes the normalized header (`Accept-Encoding: gzip`) in the corresponding origin request. The origin request doesn't include other values that were in the `Accept-Encoding` header sent by the viewer.

This behavior is the same when the viewer supports both Gzip and Brotli (the `Accept-Encoding` header in the viewer request includes both `gzip` and `br` as values), because in this scenario, caching compressed objects for Brotli is not enabled.

To understand what CloudFront does if caching compressed objects is enabled for Brotli but not Gzip, replace the two compression formats with each other in the preceding example.

If the viewer does not support the compression format for which caching is enabled (the `Accept-Encoding` header in the viewer request doesn't contain the value for that format), CloudFront:

- Doesn't include the `Accept-Encoding` header in the cache key.
- Includes `Accept-Encoding: identity` in the corresponding origin request. The origin request doesn't include other values that were in the `Accept-Encoding` header sent by the viewer.

When caching compressed objects is disabled for both compression formats

When caching compressed objects is disabled for both compression formats, CloudFront treats the `Accept-Encoding` header the same as any other HTTP header in the viewer request. By default, it's not included in the cache key and it's not included in origin requests. You can include it in the headers whitelist in a cache policy or an origin request policy like any other HTTP header.

Creating cache policies

You can use a cache policy to improve your cache hit ratio by controlling the values (URL query strings, HTTP headers, and cookies) that are included in the cache key. You can create a cache policy in the CloudFront console, with the AWS Command Line Interface (AWS CLI), or with the CloudFront API.

After you create a cache policy, you attach it to one or more cache behaviors in a CloudFront distribution.

Creating cache policies (console)

To create a cache policy (console)

1. Sign in to the AWS Management Console and open the **Policies** page in the CloudFront console at <https://console.aws.amazon.com/cloudfront/v2/home?#/policies>.
2. Choose **Create cache policy**.
3. Choose the desired setting for this cache policy. For more information, see [Understanding cache policies \(p. 84\)](#).
4. When finished, choose **Create cache policy**.

After you create a cache policy, you can attach it to a cache behavior.

To attach a cache policy to an existing distribution (console)

1. Open the **Distributions** page in the CloudFront console at <https://console.aws.amazon.com/cloudfront/home#distributions>.
2. Choose the distribution to update, then choose the **Behaviors** tab.
3. Choose the cache behavior to update, then choose **Edit**.

Or, to create a new cache behavior, choose **Create Behavior**.
4. For **Cache and origin request settings**, make sure that **Use a cache policy and origin request policy** is chosen.
5. For **Cache Policy**, choose the cache policy to attach to this cache behavior.
6. At the bottom of the page, choose **Yes, Edit**.

To attach a cache policy to a new distribution (console)

1. Open the CloudFront console at <https://console.aws.amazon.com/cloudfront/home>.
2. Choose **Create Distribution**, then for **Web**, choose **Get Started**.
3. For **Cache and origin request settings**, make sure that **Use a cache policy and origin request policy** is chosen.
4. For **Cache Policy**, choose the cache policy to attach to this distribution's default cache behavior.
5. Choose the desired settings for the origin, default cache behavior, and distribution. For more information, see [Values That You Specify When You Create or Update a Distribution \(p. 38\)](#).
6. When finished, choose **Create Distribution**.

Creating cache policies (AWS CLI)

To create a cache policy with the AWS Command Line Interface (AWS CLI), use the **aws cloudfront create-cache-policy** command. You can use an input file to provide the command's input parameters, rather than specifying each individual parameter as command line input.

To create a cache policy (CLI with input file)

1. Use the following command to create a file named `cache-policy.yaml` that contains all of the input parameters for the **create-cache-policy** command.

```
aws cloudfront create-cache-policy --generate-cli-skeleton yaml-input > cache-policy.yaml
```

Note

The `yaml-input` option is available only in [version 2 of the AWS CLI](#). With version 1 of the AWS CLI, you can generate an input file in JSON format. For more information, see [Generating AWS CLI skeleton and input parameters from a JSON or YAML input file](#) in the *AWS Command Line Interface User Guide*.

2. Open the file named `cache-policy.yaml` that you just created. Edit the file to specify the cache policy settings that you want, then save the file. You can remove optional fields from the file, but don't remove the required fields.

For more information about the cache policy settings, see [Understanding cache policies \(p. 84\)](#).

3. Use the following command to create the cache policy using input parameters from the `cache-policy.yaml` file.

```
aws cloudfront create-cache-policy --cli-input-yaml file://cache-policy.yaml
```

Make note of the `Id` value in the command's output. This is the cache policy ID, and you need it to attach the cache policy to a CloudFront distribution's cache behavior.

To attach a cache policy to an existing distribution (CLI with input file)

1. Use the following command to save the distribution configuration for the CloudFront distribution that you want to update. Replace `distribution_ID` with the distribution's ID.

```
aws cloudfront get-distribution-config --id distribution_ID --output yaml > dist-config.yaml
```

Note

The `--output yaml` option is available only in [version 2 of the AWS CLI](#). With version 1 of the AWS CLI, you can generate the output in JSON format. For more information, see [Controlling command output from the AWS CLI](#) in the *AWS Command Line Interface User Guide*.

2. Open the file named `dist-config.yaml` that you just created. Edit the file, making the following changes to each cache behavior that you are updating to use a cache policy.
 - In the cache behavior, add a field named `CachePolicyId`. For the field's value, use the cache policy ID that you noted after creating the policy.
 - Remove the `MinTTL`, `MaxTTL`, `DefaultTTL`, and `ForwardedValues` fields from the cache behavior. These settings are specified in the cache policy, so you can't include these fields and a cache policy in the same cache behavior.
 - Rename the `ETag` field to `IfMatch`, but don't change the field's value.

Save the file when finished.

3. Use the following command to update the distribution to use the cache policy. Replace `distribution_ID` with the distribution's ID.

```
aws cloudfront update-distribution --id distribution_ID --cli-input-yaml file://dist-config.yaml
```

To attach a cache policy to a new distribution (CLI with input file)

1. Use the following command to create a file named `distribution.yaml` that contains all of the input parameters for the **create-distribution** command.

```
aws cloudfront create-distribution --generate-cli-skeleton yaml-input >
distribution.yaml
```

Note

The `yaml-input` option is available only in [version 2 of the AWS CLI](#). With version 1 of the AWS CLI, you can generate an input file in JSON format. For more information, see [Generating AWS CLI skeleton and input parameters from a JSON or YAML input file](#) in the *AWS Command Line Interface User Guide*.

2. Open the file named `distribution.yaml` that you just created. In the default cache behavior, in the `CachePolicyId` field, enter the cache policy ID that you noted after creating the policy. Continue editing the file to specify the distribution settings that you want, then save the file when finished.

For more information about the distribution settings, see [Values That You Specify When You Create or Update a Distribution](#) (p. 38).

3. Use the following command to create the distribution using input parameters from the `distribution.yaml` file.

```
aws cloudfront create-distribution --cli-input-yaml file://distribution.yaml
```

Creating cache policies (API)

To create a cache policy with the CloudFront API, use [CreateCachePolicy](#). For more information about the fields that you specify in this API call, see [Understanding cache policies](#) (p. 84) and the API reference documentation for your AWS SDK or other API client.

After you create a cache policy, you can attach it to a cache behavior, using one of the following API calls:

- To attach it to a cache behavior in an existing distribution, use [UpdateDistribution](#).
- To attach it to a cache behavior in a new distribution, use [CreateDistribution](#).

For both of these API calls, provide the cache policy's ID in the `CachePolicyId` field, inside a cache behavior. For more information about the other fields that you specify in these API calls, see [Values That You Specify When You Create or Update a Distribution](#) (p. 38) and the API reference documentation for your AWS SDK or other API client.

Using the managed cache policies

CloudFront provides a set of managed cache policies that you can attach to any of your distribution's cache behaviors. With a managed cache policy, you don't need to write or maintain your own cache policy. The managed policies use settings that are optimized for specific use cases.

Contents

- [Attaching a managed cache policy](#) (p. 92)
- [Understanding the managed cache policies](#) (p. 92)

Attaching a managed cache policy

To use a managed cache policy, you attach it to a cache behavior in your distribution. The process is the same as when you create a cache policy, but instead of creating a new one, you just attach one of the managed cache policies. You attach the policy either by name (with the console) or by ID (with the AWS CLI or SDKs). The names and IDs are listed in the following section.

For more information, see [Creating cache policies \(p. 88\)](#).

Understanding the managed cache policies

The following list describes the managed cache policies.

Name: Managed-CachingOptimized, ID: 658327ea-f89d-4fab-a63d-7e88639e58f6

This policy is designed to optimize cache efficiency by minimizing the values that CloudFront includes in the cache key. CloudFront doesn't include any query strings or cookies in the cache key, and only includes the normalized `Accept-Encoding` header. This enables CloudFront to separately cache objects in the Gzip and Brotli compressions formats when the origin returns them or when [CloudFront edge compression \(p. 116\)](#) is enabled.

Policy settings

- **MinTTL:** 1 second.
- **MaxTTL:** 31,536,000 seconds (365 days).
- **DefaultTTL:** 86,400 seconds (24 hours).
- **Query strings included in cache key:** None.
- **Headers included in cache key:** None are explicitly included. The normalized `Accept-Encoding` header is included because the cache compressed objects setting is enabled. For more information, see [Cache compressed objects \(uses the Accept-Encoding header\) \(p. 86\)](#).
- **Cookies included in cache key:** None.
- **Cache compressed objects setting:** Enabled. For more information, see [Cache compressed objects \(uses the Accept-Encoding header\) \(p. 86\)](#).

Name: Managed-CachingOptimizedForUncompressedObjects, ID: b2884449-e4de-46a7-ac36-70bc7f1ddd6d

This policy is designed to optimize cache efficiency by minimizing the values included in the cache key. No query strings, headers, or cookies are included. This policy is identical to the previous one, but it disables the cache compressed objects setting.

Policy settings

- **MinTTL:** 1 second
- **MaxTTL:** 31,536,000 seconds (365 days)
- **DefaultTTL:** 86,400 seconds (24 hours)
- **Query strings included in cache key:** None
- **Headers included in cache key:** None
- **Cookies included in cache key:** None
- **Cache compressed objects setting:** Disabled

Name: Managed-CachingDisabled, ID: 4135ea2d-6df8-44a3-9df3-4b5a84be39ad

This policy disables caching. This policy is useful for dynamic content and for requests that are not cacheable.

Policy settings

- **MinTTL:** 0 seconds
- **MaxTTL:** 0 seconds
- **DefaultTTL:** 0 seconds
- **Query strings included in cache key:** None
- **Headers included in cache key:** None
- **Cookies included in cache key:** None
- **Cache compressed objects setting:** Disabled

Name: Managed-Elemental-MediaPackage, **ID:** 08627262-05a9-4f76-9ded-b50ca2e3a84f

This policy is designed for use with an origin that is an AWS Elemental MediaPackage endpoint.

Policy settings

- **MinTTL:** 0 seconds
- **MaxTTL:** 31,536,000 seconds (365 days)
- **DefaultTTL:** 86,400 seconds (24 hours)
- **Query strings included in cache key:**
 - m
 - start
 - end
 - aws.manifestfilter
- **Headers included in cache key:**
 - Origin
- **Cookies included in cache key:** None
- **Cache compressed objects setting:** Enabled. For more information, see [Cache compressed objects \(uses the Accept-Encoding header\)](#) (p. 86).

Understanding the cache key

The *cache key* determines whether a viewer request to a CloudFront edge location results in a *cache hit*. The cache key is the unique identifier for an object in the cache. Each object in the cache has a unique cache key.

A cache hit occurs when a viewer request generates the same cache key as a prior request, and the object for that cache key is in the edge location's cache and valid. When there's a cache hit, the requested object is served to the viewer from a CloudFront edge location, which has the following benefits:

- Reduced load on your origin server
- Reduced latency for the viewer

You can get better performance from your website or application when you have a higher *cache hit ratio* (a higher proportion of viewer requests that result in a cache hit). One way to improve your cache hit ratio is to include only the minimum necessary values in the cache key. For more information, see the following sections.

You can modify the values (URL query strings, HTTP headers, and cookies) in the cache key by using a [cache policy](#) (p. 83). (You can also modify the cache key using a [Lambda@Edge function](#) (p. 318).) Before modifying the cache key, it's important to understand how your application is designed and when and how it might serve different responses based on characteristics of the viewer request. When a value in the viewer request determines the response that your origin returns, you should include that value in the cache key. But if you include a value in the cache key that doesn't affect the response that your origin returns, you might end up caching duplicate objects.

The default cache key

By default, the cache key for a CloudFront distribution includes the following information:

- The domain name of the CloudFront distribution (for example, `d111111abcdef8.cloudfront.net`)
- The URL path of the requested object (for example, `/content/stories/example-story.html`)

Other values from the viewer request are not included in the cache key, by default. Consider the following HTTP request from a web browser.

```
GET /content/stories/example-story.html?ref=0123abc&split-pages=false HTTP/1.1
Host: d111111abcdef8.cloudfront.net
User-Agent: Mozilla/5.0 Gecko/20100101 Firefox/68.0
Accept: text/html, */*
Accept-Language: en-US,en
Cookie: session_id=01234abcd
Referer: https://news.example.com/
```

When a viewer request like this one comes in to a CloudFront edge location, CloudFront uses the cache key to determine if there's a cache hit. By default, only the information shown in bold is included in the cache key. If the requested object is not in the cache (a cache miss), then CloudFront sends a request to the origin to get the object. After getting the object, CloudFront returns it to the viewer and stores it in the edge location's cache.

When CloudFront receives another request for the same object, as determined by the cache key, CloudFront serves the cached object to the viewer immediately, without sending a request to the origin. For example, consider the following HTTP request that comes in after the previous request.

```
GET /content/stories/example-story.html?ref=xyz987&split-pages=true HTTP/1.1
Host: d111111abcdef8.cloudfront.net
User-Agent: Mozilla/5.0 AppleWebKit/537.36 Chrome/83.0.4103.116
Accept: text/html, */*
Accept-Language: en-US,en
Cookie: session_id=wxyz9876
Referer: https://rss.news.example.net/
```

This request is for the same object as the previous request, but is different from the previous request. It has a different URL query string, different `User-Agent` and `Referer` headers, and a different `session_id` cookie. However, none of these values are part of the cache key by default, so this second request results in a cache hit.

Customizing the cache key

In some cases, you might want to include more information in the cache key, even though doing so might result in fewer cache hits. You specify what to include in the cache key by using a [cache policy](#) (p. 83).

For example, if your origin server uses the `Accept-Language` HTTP header in viewer requests to return different content based on the viewer's language, you might want to include this header in the cache key. When you do that, CloudFront uses this header to determine cache hits, and includes the header in *origin requests* (requests that CloudFront sends to the origin when there's a cache miss).

One potential consequence of including additional values in the cache key is that CloudFront might end up caching duplicate objects because of the variation that can occur in viewer requests. For example, viewers might send any of the following values for the `Accept-Language` header:

- en-US, en
- en, en-US
- en-US, en
- en-US

All of these different values indicate that the viewer's language is English, but the variation can cause CloudFront to cache the same object multiple times. This can reduce cache hits and increase the number of origin requests. You could avoid this duplication by not including the `Accept-Language` header in the cache key, and instead configuring your website or application to use different URLs for content in different languages (for example, `/en-US/content/stories/example-story.html`).

For any given value that you intend to include in the cache key, you should make sure that you understand how many different variations of that value might appear in viewer requests. For certain request values, it rarely makes sense to include them in the cache key. For example, the `User-Agent` header can have thousands of unique variations, so it's generally not a good candidate for including in the cache key. Cookies that have user-specific or session-specific values and are unique across thousands (or even millions) of requests are also not good candidates for cache key inclusion. If you do include these values in the cache key, each unique variation results in another copy of the object in the cache. If these copies of the object are not unique, or if you end up with such a large number of slightly different objects that each object only gets a small number of cache hits, you might want to consider a different approach. You can exclude these highly variable values from the cache key, or you can mark objects as non-cacheable.

Use caution when customizing the cache key. Sometimes it's desirable, but it can have unintended consequences such as caching duplicate objects, lowering your cache hit ratio, and increasing the number of origin requests. If your origin website or application needs to receive certain values from viewer requests for analytics, telemetry, or other uses, but these values don't change the object that the origin returns, use an [origin request policy \(p. 95\)](#) to include these values in origin requests but *not* include them in the cache key.

Controlling origin requests

When a viewer request to CloudFront results in a *cache miss* (the requested object is not cached at the edge location), CloudFront sends a request to the origin to retrieve the object. This is called an *origin request*. The origin request always includes the following information from the viewer request:

- The URL path (the path only, without URL query strings or the domain name)
- The request body (if there is one)
- The HTTP headers that CloudFront automatically includes in every origin request, including `Host`, `User-Agent`, and `X-Amz-Cf-Id`.

Other information from the viewer request, such as URL query strings, HTTP headers, and cookies, is not included in the origin request by default. But you might want to receive some of this other information at the origin, for example to collect data for analytics or telemetry. You can use an *origin request policy* to control the information that's included in an origin request.

Origin request policies are separate from [cache policies \(p. 83\)](#), which control the cache key. This separation enables you to receive additional information at the origin and also maintain a good *cache hit ratio* (the proportion of viewer requests that result in a cache hit). You do this by separately controlling which information is included in origin requests (using the origin request policy) and which is included in the cache key (using the cache policy).

Although the two kinds of policies are separate, they are related. All URL query strings, HTTP headers, and cookies that you include in the cache key (using a cache policy) are automatically included in origin

requests. Use the origin request policy to specify the information that you want to include in origin requests, but *not* include in the cache key. Just like a cache policy, you attach an origin request policy to one or more cache behaviors in a CloudFront distribution.

You can also use an origin request policy to add additional HTTP headers to an origin request that were not included in the viewer request. These additional headers are added by CloudFront before sending the origin request, with header values that are determined automatically based on the viewer request. For more information, see [Using the CloudFront HTTP headers \(p. 102\)](#).

Contents

- [Understanding origin request policies \(p. 96\)](#)
 - [Policy information \(p. 96\)](#)
 - [Origin request settings \(p. 96\)](#)
- [Creating origin request policies \(p. 97\)](#)
- [Using the managed origin request policies \(p. 100\)](#)
 - [Attaching a managed origin request policy \(p. 101\)](#)
 - [Understanding the managed origin request policies \(p. 101\)](#)

Understanding origin request policies

CloudFront provides some predefined origin request policies, known as *managed policies*, for common use cases. You can use these managed policies, or you can create your own origin request policy that's specific to your needs. For more information about the managed policies, see [Using the managed origin request policies \(p. 100\)](#).

An origin request policy contains the following settings, which are categorized into *policy information* and *origin request settings*.

Policy information

Name

A name to identify the origin request policy. In the console, you use the name to attach the origin request policy to a cache behavior.

Comment

A comment to describe the origin request policy. This is optional.

Origin request settings

Origin request settings specify the values in viewer requests that are included in requests that CloudFront sends to the origin (known as origin requests). The values can include URL query strings, HTTP headers, and cookies. The values that you specify are included in origin requests, but are not included in the cache key. For information about controlling the cache key, see [Controlling the cache key \(p. 83\)](#).

Query strings

The URL query strings in viewer requests that CloudFront includes in origin requests. For query strings, you can choose one of the following settings:

- **None** – The query strings in viewer requests are *not* included in origin requests.
- **All** – All query strings in viewer requests are included in origin requests.
- **Whitelist** – You specify which of the query strings in viewer requests are included in origin requests.

When you use the **Whitelist** setting, you specify query strings by their name, not their value. For example, consider the following URL path:

```
/content/stories/example-story.html?split-pages=false
```

In this case, you specify the query string as `split-pages`, not as `split-pages=false`. However, CloudFront includes the full query string, including its value, in origin requests.

Headers

The HTTP headers in viewer requests that CloudFront includes in origin requests. For headers, you can choose one of the following settings:

- **None** – The HTTP headers in viewer requests are *not* included in origin requests.
- **All viewer headers** – All HTTP headers in viewer requests are included in origin requests.
- **Whitelist** – You specify which HTTP headers are included in origin requests.
- **All viewer headers and whitelisted CloudFront-* headers** – All HTTP headers in viewer requests are included in origin requests. Additionally, you specify which of the CloudFront headers you want to add to origin requests. For more information about the CloudFront headers, see [Using the CloudFront HTTP headers \(p. 102\)](#).

When you use the **Whitelist** or **All viewer headers and whitelisted CloudFront-* headers** setting, you specify HTTP headers by their name, not their value. For example, consider the following HTTP header:

```
Accept-Language: en-US,en;q=0.5
```

In this case, you specify the header as `Accept-Language`, not as `Accept-Language: en-US,en;q=0.5`. However, CloudFront includes the full header, including its value, in origin requests.

Cookies

The cookies in viewer requests that CloudFront includes in origin requests. For cookies, you can choose one of the following settings:

- **None** – The cookies in viewer requests are *not* included in origin requests.
- **All** – All cookies in viewer requests are included in origin requests.
- **Whitelist** – You specify which of the cookies in viewer requests are included in origin requests.

When you use the **Whitelist** setting, you specify cookies by their name, not their value. For example, consider the following `Cookie` header:

```
Cookie: session_ID=abcd1234
```

In this case, you specify the cookie as `session_ID`, not as `session_ID=abcd1234`. However, CloudFront includes the full cookie, including its value, in origin requests.

Creating origin request policies

You can use an origin request policy to control the values (URL query strings, HTTP headers, and cookies) that are included in requests that CloudFront sends to your origin. You can create an origin request policy in the CloudFront console, with the AWS Command Line Interface (AWS CLI), or with the CloudFront API.

After you create an origin request policy, you attach it to one or more cache behaviors in a CloudFront distribution.

Origin request policies are not required. When a cache behavior does not have an origin request policy attached, the origin request includes all the values that are specified in the [cache policy \(p. 84\)](#), but nothing more.

Note

To use an origin request policy, the cache behavior must also use a [cache policy \(p. 83\)](#). You cannot use an origin request policy in a cache behavior without a cache policy.

Creating origin request policies (console)

To create an origin request policy (console)

1. Sign in to the AWS Management Console and open the **Policies** page in the CloudFront console at <https://console.aws.amazon.com/cloudfront/v2/home?#/policies>.
2. Choose **Origin request policy**, then choose **Create origin request policy**.
3. Choose the desired setting for this origin request policy. For more information, see [Understanding origin request policies \(p. 96\)](#).
4. When finished, choose **Create origin request policy**.

After you create an origin request policy, you can attach it to a cache behavior.

To attach an origin request policy to an existing distribution (console)

1. Open the **Distributions** page in the CloudFront console at <https://console.aws.amazon.com/cloudfront/home#distributions>.
2. Choose the distribution to update, then choose the **Behaviors** tab.
3. Choose the cache behavior to update, then choose **Edit**.

Or, to create a new cache behavior, choose **Create Behavior**.

4. For **Cache and origin request settings**, make sure that **Use a cache policy and origin request policy** is chosen.
5. For **Origin Request Policy**, choose the origin request policy to attach to this cache behavior.
6. At the bottom of the page, choose **Yes, Edit**.

To attach an origin request policy to a new distribution (console)

1. Open the CloudFront console at <https://console.aws.amazon.com/cloudfront/home>.
2. Choose **Create Distribution**, then for **Web**, choose **Get Started**.
3. For **Cache and origin request settings**, make sure that **Use a cache policy and origin request policy** is chosen.
4. For **Origin Request Policy**, choose the origin request policy to attach to this distribution's default cache behavior.
5. Choose the desired settings for the origin, default cache behavior, and distribution. For more information, see [Values That You Specify When You Create or Update a Distribution \(p. 38\)](#).
6. When finished, choose **Create Distribution**.

Creating origin request policies (AWS CLI)

To create an origin request policy with the AWS Command Line Interface (AWS CLI), use the **aws cloudfront create-origin-request-policy** command. You can use an input file to provide the command's input parameters, rather than specifying each individual parameter as command line input.

To create an origin request policy (CLI with input file)

1. Use the following command to create a file named `origin-request-policy.yaml` that contains all of the input parameters for the **create-origin-request-policy** command.

```
aws cloudfront create-origin-request-policy --generate-cli-skeleton yaml-input >
origin-request-policy.yaml
```

Note

The `yaml-input` option is available only in [version 2 of the AWS CLI](#). With version 1 of the AWS CLI, you can generate an input file in JSON format. For more information, see [Generating AWS CLI skeleton and input parameters from a JSON or YAML input file](#) in the *AWS Command Line Interface User Guide*.

2. Open the file named `origin-request-policy.yaml` that you just created. Edit the file to specify the origin request policy settings that you want, then save the file. You can remove optional fields from the file, but don't remove the required fields.

For more information about the origin request policy settings, see [Understanding origin request policies \(p. 96\)](#).

3. Use the following command to create the origin request policy using input parameters from the `origin-request-policy.yaml` file.

```
aws cloudfront create-origin-request-policy --cli-input-yaml file://origin-request-
policy.yaml
```

Make note of the `Id` value in the command's output. This is the origin request policy ID, and you need it to attach the origin request policy to a CloudFront distribution's cache behavior.

To attach an origin request policy to an existing distribution (CLI with input file)

1. Use the following command to save the distribution configuration for the CloudFront distribution that you want to update. Replace `distribution_ID` with the distribution's ID.

```
aws cloudfront get-distribution-config --id distribution_ID --output yaml > dist-
config.yaml
```

Note

The `--output yaml` option is available only in [version 2 of the AWS CLI](#). With version 1 of the AWS CLI, you can generate the output in JSON format. For more information, see [Controlling command output from the AWS CLI](#) in the *AWS Command Line Interface User Guide*.

2. Open the file named `dist-config.yaml` that you just created. Edit the file, making the following changes to each cache behavior that you are updating to use an origin request policy.
 - In the cache behavior, add a field named `OriginRequestPolicyId`. For the field's value, use the origin request policy ID that you noted after creating the policy.
 - Rename the `ETag` field to `IfMatch`, but don't change the field's value.

Save the file when finished.

3. Use the following command to update the distribution to use the origin request policy. Replace `distribution_ID` with the distribution's ID.

```
aws cloudfront update-distribution --id distribution_ID --cli-input-yaml file://dist-config.yaml
```

To attach a cache policy to a new distribution (CLI with input file)

1. Use the following command to create a file named `distribution.yaml` that contains all of the input parameters for the **create-distribution** command.

```
aws cloudfront create-distribution --generate-cli-skeleton yaml-input > distribution.yaml
```

Note

The `yaml-input` option is available only in [version 2 of the AWS CLI](#). With version 1 of the AWS CLI, you can generate an input file in JSON format. For more information, see [Generating AWS CLI skeleton and input parameters from a JSON or YAML input file](#) in the *AWS Command Line Interface User Guide*.

2. Open the file named `distribution.yaml` that you just created. In the default cache behavior, in the `OriginRequestPolicyId` field, enter the origin request policy ID that you noted after creating the policy. Continue editing the file to specify the distribution settings that you want, then save the file when finished.

For more information about the distribution settings, see [Values That You Specify When You Create or Update a Distribution \(p. 38\)](#).

3. Use the following command to create the distribution using input parameters from the `distribution.yaml` file.

```
aws cloudfront create-distribution --cli-input-yaml file://distribution.yaml
```

Creating origin request policies (API)

To create an origin request policy with the CloudFront API, use [CreateOriginRequestPolicy](#). For more information about the fields that you specify in this API call, see [Understanding origin request policies \(p. 96\)](#) and the API reference documentation for your AWS SDK or other API client.

After you create an origin request policy, you can attach it to a cache behavior, using one of the following API calls:

- To attach it to a cache behavior in an existing distribution, use [UpdateDistribution](#).
- To attach it to a cache behavior in a new distribution, use [CreateDistribution](#).

For both of these API calls, provide the origin request policy's ID in the `OriginRequestPolicyId` field, inside a cache behavior. For more information about the other fields that you specify in these API calls, see [Values That You Specify When You Create or Update a Distribution \(p. 38\)](#) and the API reference documentation for your AWS SDK or other API client.

Using the managed origin request policies

CloudFront provides a set of managed origin request policies that you can attach to any of your distribution's cache behaviors. With a managed origin request policy, you don't need to write or maintain

your own origin request policy. The managed policies use settings that are optimized for specific use cases.

Contents

- [Attaching a managed origin request policy \(p. 101\)](#)
- [Understanding the managed origin request policies \(p. 101\)](#)

Attaching a managed origin request policy

To use a managed origin request policy, you attach it to a cache behavior in your distribution. The process is the same as when you create an origin request policy, but instead of creating a new one, you just attach one of the managed origin request policies. You attach the policy either by name (with the console) or by ID (with the AWS CLI or SDKs). The names and IDs are listed in the following section.

For more information, see [Creating origin request policies \(p. 97\)](#).

Understanding the managed origin request policies

The following list describes the managed origin request policies.

Name: Managed-UserAgentRefererHeaders, ID: acba4595-bd28-49b8-b9fe-13317c0390fa

This policy includes only the `User-Agent` and `Referer` headers. It doesn't include any query strings or cookies.

Policy settings

- **Query strings included in origin requests:** None
- **Headers included in origin requests:**
 - `User-Agent`
 - `Referer`
- **Cookies included in origin requests:** None

Name: Managed-CORS-CustomOrigin, ID: 59781a5b-3903-41f3-afcb-af62929ccde1

This policy includes the header that enables cross-origin resource sharing (CORS) requests when the origin is a custom origin.

Policy settings

- **Query strings included in origin requests:** None
- **Headers included in origin requests:**
 - `Origin`
- **Cookies included in origin requests:** None

Name: Managed-CORS-S3Origin, ID: 88a5eaf4-2fd4-4709-b370-b4c650ea3fcf

This policy includes the headers that enable cross-origin resource sharing (CORS) requests when the origin is an Amazon S3 bucket.

Policy settings

- **Query strings included in origin requests:** None
- **Headers included in origin requests:**
 - `Origin`
 - `Access-Control-Request-Headers`
 - `Access-Control-Request-Method`
- **Cookies included in origin requests:** None

Name: Managed-AllViewer, ID: 216adef6-5c7f-47e4-b989-5492eafa07d3

This policy includes all values (query strings, headers, and cookies) in the viewer request.

Policy settings

- **Query strings included in origin requests:** All
- **Headers included in origin requests:** All headers in the viewer request
- **Cookies included in origin requests:** All

Name: Managed-Elemental-MediaTailor-PersonalizedManifests, ID: 775133bc-15f2-49f9-abea-afb2e0bf67d2

This policy is designed for use with an origin that is an AWS Elemental MediaTailor endpoint.

Policy settings

- **Query strings included in origin requests:** All
- **Headers included in origin requests:**
 - Origin
 - Access-Control-Request-Headers
 - Access-Control-Request-Method
 - User-Agent
 - X-Forwarded-For
- **Cookies included in origin requests:** None

Using the CloudFront HTTP headers

You can configure CloudFront to add specific HTTP headers based on characteristics of the viewer request. With these headers, your origin can receive information about the viewer's device type, geographic location, and more, without the need for custom code to determine this information. If your origin returns different responses based on the information in these headers, you can include them in the *cache key* so that CloudFront caches the different responses separately.

To receive these headers at your origin, use an *origin request policy*. For more information, see [Controlling origin requests](#) (p. 95).

To include these headers in the cache key, use a *cache policy*. For more information, see [Controlling the cache key](#) (p. 83) and [Understanding the cache key](#) (p. 93).

Topics

- [Headers for determining the viewer's device type](#) (p. 102)
- [Headers for determining the viewer's location](#) (p. 103)
- [Other CloudFront headers](#) (p. 103)

Headers for determining the viewer's device type

Use the following headers to determine the viewer's device type. Based on the value of the `User-Agent` header, CloudFront sets the value of these headers to `true` or `false`. If a device falls into more than one category, more than one value might be `true`. For example, for some tablet devices, CloudFront might set both `CloudFront-Is-Mobile-Viewer` and `CloudFront-Is-Tablet-Viewer` to `true`.

- `CloudFront-Is-Android-Viewer` – Set to `true` when CloudFront determines that the viewer is a device with the Android operating system.

- `CloudFront-Is-Desktop-Viewer` – Set to `true` when CloudFront determines that the viewer is a desktop device.
- `CloudFront-Is-IOS-Viewer` – Set to `true` when CloudFront determines that the viewer is a device with an Apple operating system.
- `CloudFront-Is-Mobile-Viewer` – Set to `true` when CloudFront determines that the viewer is a mobile device.
- `CloudFront-Is-SmartTV-Viewer` – Set to `true` when CloudFront determines that the viewer is a smart TV.
- `CloudFront-Is-Tablet-Viewer` – Set to `true` when CloudFront determines that the viewer is a tablet.

Headers for determining the viewer's location

Use the following headers to determine the viewer's location. CloudFront determines the values for these headers based on the viewer's IP address.

Note

For non-ASCII characters in these headers' values, CloudFront percent encodes the character according to [section 1.2 of RFC 3986](#).

- `CloudFront-Viewer-City` – Contains the name of the viewer's city.
- `CloudFront-Viewer-Country` – Contains the two-letter country code for the viewer's country. For a list of country codes, see [ISO 3166-1 alpha-2](#).
- `CloudFront-Viewer-Country-Name` – Contains the name of the viewer's country.
- `CloudFront-Viewer-Country-Region` – Contains a code (up to three characters) that represent the viewer's region. The region is the most specific subdivision of the [ISO 3166-2](#) code.
- `CloudFront-Viewer-Country-Region-Name` – Contains the name of the viewer's region. The region is the most specific subdivision of the [ISO 3166-2](#) code.
- `CloudFront-Viewer-Latitude` – Contains the viewer's approximate latitude.
- `CloudFront-Viewer-Longitude` – Contains the viewer's approximate longitude.
- `CloudFront-Viewer-Metro-Code` – Contains the viewer's metro code. This is present only when the viewer is in the United States.
- `CloudFront-Viewer-Postal-Code` – Contains the viewer's postal code.
- `CloudFront-Viewer-Time-Zone` – Contains the viewer's time zone, in [IANA time zone database format](#) (for example, `America/Los_Angeles`).

Other CloudFront headers

Use the following headers to determine the protocol and HTTP version of the viewer's request.

- `CloudFront-Forwarded-Proto` – Contains the protocol of the viewer's request (HTTP or HTTPS).
- `CloudFront-Viewer-Http-Version` – Contains the HTTP version of the viewer's request.

Adding, Removing, or Replacing Content That CloudFront Distributes

This section explains how to make sure CloudFront can access the content that you want to be served to your viewers, how to specify the objects in your website or in your application, and how to remove or replace content.

Topics

- [Adding and Accessing Content That CloudFront Distributes \(p. 104\)](#)
- [Updating Existing Content with a CloudFront Distribution \(p. 104\)](#)
- [Removing Content so CloudFront Won't Distribute It \(p. 106\)](#)
- [Customizing the URL Format for Files in CloudFront \(p. 106\)](#)
- [Specifying a Default Root Object \(p. 107\)](#)
- [Invalidating Files \(p. 110\)](#)
- [Serving compressed files \(p. 116\)](#)

Adding and Accessing Content That CloudFront Distributes

When you want CloudFront to distribute content (objects), you add files to one of the origins that you specified for the distribution, and you expose a CloudFront link to the files. A CloudFront edge location doesn't fetch the new files from an origin until the edge location receives viewer requests for them. For more information, see [How CloudFront delivers content \(p. 5\)](#).

When you add a file that you want CloudFront to distribute, make sure that you add it to one of the Amazon S3 buckets specified in your distribution or, for a custom origin, to a directory in the specified domain. In addition, confirm that the path pattern in the applicable cache behavior sends requests to the correct origin.

For example, suppose the path pattern for a cache behavior is `*.html`. If you don't have any other cache behaviors configured to forward requests to that origin, CloudFront will only forward `*.html` files. In this scenario, for example, CloudFront will never distribute `.jpg` files that you upload to the origin, because you haven't created a cache behavior that includes `.jpg` files.

CloudFront servers don't determine the MIME type for the objects that they serve. When you upload a file to your origin, we recommend that you set the `Content-Type` header field for it.

Updating Existing Content with a CloudFront Distribution

There are two ways to update existing content that CloudFront is set up to distribute for you:

- Update files by using the same name
- Update by using a version identifier in the file name

We recommend that you use a version identifier in file names or in folder names, to help give you more control over managing the content that CloudFront serves.

Updating Existing Files Using Versioned File Names

When you update existing files in a CloudFront distribution, we recommend that you include some sort of version identifier either in your file names or in your directory names to give yourself better control over your content. This identifier might be a date-time stamp, a sequential number, or some other method of distinguishing two versions of the same object.

For example, instead of naming a graphic file `image.jpg`, you might call it `image_1.jpg`. When you want to start serving a new version of the file, you'd name the new file `image_2.jpg`, and you'd update the links in your web application or website to point to `image_2.jpg`. Alternatively, you might put all graphics in an `images_v1` directory and, when you want to start serving new versions of one or more graphics, you'd create a new `images_v2` directory, and you'd update your links to point to that directory. With versioning, you don't have to wait for an object to expire before CloudFront begins to serve a new version of it, and you don't have to pay for object invalidation.

Even if you version your files, we still recommend that you set an expiration date. For more information, see [Managing How Long Content Stays in an Edge Cache \(Expiration\)](#) (p. 236).

Note

Specifying versioned file names or directory names is not related to Amazon S3 object versioning.

Updating Existing Content Using the Same File Names

Although you can update existing files in a CloudFront distribution and use the same file names, we don't recommend it. CloudFront distributes files to edge locations only when the files are requested, not when you put new or updated files in your origin. If you update an existing file in your origin with a newer version that has the same name, an edge location won't get that new version from your origin until both of the following occur:

- The old version of the file in the cache expires. For more information, see [Managing How Long Content Stays in an Edge Cache \(Expiration\)](#) (p. 236).
- There's a user request for the file at that edge location.

If you use the same names when you replace files, you can't control when CloudFront starts to serve the new files. By default, CloudFront caches files in edge locations for 24 hours. (For more information, see [Managing How Long Content Stays in an Edge Cache \(Expiration\)](#) (p. 236).) For example, if you're replacing all of the files on an entire website:

- Files for the less popular pages may not be in any edge locations. The new versions of these files will start being served on the next request.
- Files for some pages may be in some edge locations and not in others, so your end users will see different versions depending on which edge location they're served from.
- New versions of the files for the most popular pages might not be served for up to 24 hours because CloudFront might have retrieved the files for those pages just before you replaced the files with new versions.

Removing Content so CloudFront Won't Distribute It

You can remove files from your origin that you no longer want to be included in your CloudFront distribution. However, CloudFront will continue to show viewers content from the edge cache until the files expire.

If you want to remove a file right away, you must do one of the following:

- **Invalidate the file.** For more information, see [Invalidating Files \(p. 110\)](#).
- **Use file versioning.** When you use versioning, different versions of a file have different names that you can use in your CloudFront distribution, to change which file is returned to viewers. For more information, see [Updating Existing Files Using Versioned File Names \(p. 105\)](#).

Customizing the URL Format for Files in CloudFront

After you set up your origin with the objects (content) that you want CloudFront to serve to your viewers, you must use the correct URLs to reference those objects in your website or application code so that CloudFront can serve it.

The domain name that you use in the URLs for objects on your web pages or in your web application can be either of the following:

- The domain name, such as `d111111abcdef8.cloudfront.net`, that CloudFront automatically assigns when you create a distribution
- Your own domain name, such as `example.com`

For example, you might use one of the following URLs to return the file `image.jpg`:

```
http://d111111abcdef8.cloudfront.net/images/image.jpg
```

```
http://example.com/images/image.jpg
```

You use the same URL format whether you store the content in Amazon S3 buckets or at a custom origin, like one of your own web servers.

Note

The URL format depends in part on the value that you specify for **Origin Path** in your distribution. This value gives CloudFront a top directory path for your objects. For more information about setting the origin path when you create a web distribution, see [Origin Path \(p. 42\)](#).

For more information about URL formats, see the following sections.

Using Your Own Domain Name (Example.com)

Instead of using the default domain name that CloudFront assigns for you when you create a distribution, you can [add an alternate domain name](#) that's easier to work with, like `example.com`. By setting up your own domain name with CloudFront, you can use a URL like this for objects in your distribution:

```
http://example.com/images/image.jpg
```

If you plan to use HTTPS between viewers and CloudFront, see [Using Alternate Domain Names and HTTPS \(p. 132\)](#).

Using a Trailing Slash (/) in URLs

When you specify URLs for directories in your CloudFront distribution, choose either to always use a trailing slash or to never use a trailing slash. For example, choose only one of the following formats for all of your URLs:

```
http://d1111111abcdef8.cloudfront.net/images/
```

```
http://d1111111abcdef8.cloudfront.net/images
```

Why does it matter?

Both formats work to link to CloudFront objects, but being consistent can help prevent issues when you want to invalidate a directory later. CloudFront stores URLs exactly as they are defined, including trailing slashes. So if your format is inconsistent, you'll need to invalidate directory URLs with and without the slash, to ensure that CloudFront removes the directory.

It's inconvenient to have to invalidate both URL formats, and it can lead to additional costs. That's because if you must double up invalidations to cover both types of URLs, you might exceed the maximum number of free invalidations allowed for the month. And if that happens, you'll have to pay for all the invalidations, even if only one format for each directory URL exists in CloudFront.

Creating Signed URLs for Restricted Content

If you have content that you want to restrict access to, you can create signed URLs. For example, if you want to distribute your content only to users who have authenticated, you can create URLs that are valid only for a specified time period or that are available only from a specified IP address. For more information, see [Serving Private Content with Signed URLs and Signed Cookies \(p. 145\)](#).

Creating URLs for RTMP Distribution Media Files

If you use an RTMP distribution to serve streaming content, you must include additional characters in the URLs for your files. For more information, see [Configuring the Media Player \(p. 314\)](#).

Specifying a Default Root Object

You can configure CloudFront to return a specific object (the default root object) when a user requests the root URL for your web distribution instead of requesting an object in your distribution. Specifying a default root object lets you avoid exposing the contents of your distribution.

Topics

- [How to Specify a Default Root Object \(p. 107\)](#)
- [How Headers Work With Default Root Objects \(p. 108\)](#)
- [How CloudFront Works if You Don't Define a Root Object \(p. 109\)](#)

How to Specify a Default Root Object

To avoid exposing the contents of your web distribution or returning an error, specify a default root object for your distribution by completing the following steps.

To specify a default root object for your distribution

1. Upload the default root object to the origin that your distribution points to.

The file can be any type supported by CloudFront. For a list of constraints on the file name, see the description of the `DefaultRootObject` element in [DistributionConfig Complex Type](#).

Note

If the file name of the default root object is too long or contains an invalid character, CloudFront returns the error `HTTP 400 Bad Request - InvalidDefaultRootObject`. In addition, CloudFront caches the code for one minute and writes the results to the access logs.

2. Confirm that the permissions for the object grant CloudFront at least `read` access.

For more information about Amazon S3 permissions, see [Access Control](#) in the *Amazon Simple Storage Service Developer Guide*. For information on using the Amazon S3 console to update permissions, go to the [Amazon Simple Storage Service Console User Guide](#).

3. Update your distribution to refer to the default root object using the CloudFront console or the CloudFront API.

To specify a default root object using the CloudFront console:

- a. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
- b. In the list of distributions in the top pane, select the distribution to update.
- c. In the **Distribution Details** pane, on the **General** tab, choose **Edit**.
- d. In the **Edit Distribution** dialog box, in the **Default Root Object** field, enter the file name of the default root object.

Enter only the object name, for example, `index.html`. Do not add a `/` before the object name.

- e. To save your changes, choose **Yes, Edit**.

To update your configuration using the CloudFront API, you specify a value for the `DefaultRootObject` element in your distribution. For information about using the CloudFront API to specify a default root object, see [PUT Distribution Config](#) in the *Amazon CloudFront API Reference*.

4. Confirm that you have enabled the default root object by requesting your root URL. If your browser doesn't display the default root object, perform the following steps:
 - a. Confirm that your distribution is fully deployed by viewing the status of your distribution in the CloudFront console.
 - b. Repeat steps 2 and 3 to verify that you granted the correct permissions and that you correctly updated the configuration of your distribution to specify the default root object.

How Headers Work With Default Root Objects

Here's an example of how a default root object works. Suppose the following request points to the object `image.jpg`:

```
http://d111111abcdef8.cloudfront.net/image.jpg
```

In contrast, the following request points to the root URL of the same distribution instead of to a specific object, as in the first example:

```
http://d111111abcdef8.cloudfront.net/
```

When you define a default root object, an end-user request that calls the root of your distribution returns the default root object. For example, if you designate the file `index.html` as your default root object, a request for:

```
http://d111111abcdef8.cloudfront.net/
```

Returns:

```
http://d111111abcdef8.cloudfront.net/index.html
```

However, if you define a default root object, an end-user request for a subdirectory of your distribution does not return the default root object. For example, suppose `index.html` is your default root object and that CloudFront receives an end-user request for the `install` directory under your CloudFront distribution:

```
http://d111111abcdef8.cloudfront.net/install/
```

CloudFront does not return the default root object even if a copy of `index.html` appears in the `install` directory.

If you configure your distribution to allow all of the HTTP methods that CloudFront supports, the default root object applies to all methods. For example, if your default root object is `index.php` and you write your application to submit a `POST` request to the root of your domain (`http://example.com`), CloudFront sends the request to `http://example.com/index.php`.

The behavior of CloudFront default root objects is different from the behavior of Amazon S3 index documents. When you configure an Amazon S3 bucket as a website and specify the index document, Amazon S3 returns the index document even if a user requests a subdirectory in the bucket. (A copy of the index document must appear in every subdirectory.) For more information about configuring Amazon S3 buckets as websites and about index documents, see the [Hosting Websites on Amazon S3](#) chapter in the *Amazon Simple Storage Service Developer Guide*.

Important

Remember that a default root object applies only to your CloudFront distribution. You still need to manage security for your origin. For example, if you are using an Amazon S3 origin, you still need to set your Amazon S3 bucket ACLs appropriately to ensure the level of access you want on your bucket.

How CloudFront Works if You Don't Define a Root Object

If you don't define a default root object, requests for the root of your distribution pass to your origin server. If you are using an Amazon S3 origin, any of the following might be returned:

- **A list of the contents of your Amazon S3 bucket**—Under any of the following conditions, the contents of your origin are visible to anyone who uses CloudFront to access your distribution:
 - Your bucket is not properly configured.
 - The Amazon S3 permissions on the bucket associated with your distribution and on the objects in the bucket grant access to *everyone*.
 - An end user accesses your origin using your origin root URL.
- **A list of the private contents of your origin**—If you configure your origin as a private distribution (only you and CloudFront have access), the contents of the Amazon S3 bucket associated with your distribution are visible to anyone who has the credentials to access your distribution through CloudFront. In this case, users are not able to access your content through your origin root URL. For more information about distributing private content, see [Serving Private Content with Signed URLs and Signed Cookies](#) (p. 145).

- **Error 403 Forbidden**—CloudFront returns this error if the permissions on the Amazon S3 bucket associated with your distribution or the permissions on the objects in that bucket deny access to CloudFront and to everyone.

Invalidating Files

If you need to remove a file from CloudFront edge caches before it expires, you can do one of the following:

- Invalidate the file from edge caches. The next time a viewer requests the file, CloudFront returns to the origin to fetch the latest version of the file.
- Use file versioning to serve a different version of the file that has a different name. For more information, see [Updating Existing Files Using Versioned File Names \(p. 105\)](#).

Important

You cannot invalidate objects that are served by an RTMP distribution.

To invalidate files, you can specify either the path for individual files or a path that ends with the * wildcard, which might apply to one file or to many, as shown in the following examples:

- /images/image1.jpg
- /images/image*
- /images/*

Note

If you use the AWS command line interface (CLI) for invalidating files and you specify a path that includes the * wildcard, you must use quotes (") around the path.

For example: `aws cloudfront create-invalidation --distribution-id distribution_ID --paths "/*"`

You can submit a certain number of invalidation paths each month for free. If you submit more than the allotted number of invalidation paths in a month, you pay a fee for each invalidation path that you submit. For more information about the charges for invalidation, see [Paying for File Invalidation \(p. 116\)](#).

Topics

- [Choosing Between Invalidating Files and Using Versioned File Names \(p. 110\)](#)
- [Determining Which Files to Invalidate \(p. 111\)](#)
- [Specifying the Files to Invalidate \(p. 111\)](#)
- [Invalidating Files Using the Console \(p. 113\)](#)
- [Invalidating Files Using the CloudFront API \(p. 115\)](#)
- [Concurrent Invalidation Request Maximum \(p. 116\)](#)
- [Paying for File Invalidation \(p. 116\)](#)

Choosing Between Invalidating Files and Using Versioned File Names

To control the versions of files that are served from your distribution, you can either invalidate files or give them versioned file names. If you want to update your files frequently, we recommend that you primarily use file versioning for the following reasons:

- Versioning enables you to control which file a request returns even when the user has a version cached either locally or behind a corporate caching proxy. If you invalidate the file, the user might continue to see the old version until it expires from those caches.
- CloudFront access logs include the names of your files, so versioning makes it easier to analyze the results of file changes.
- Versioning provides a way to serve different versions of files to different users.
- Versioning simplifies rolling forward and back between file revisions.
- Versioning is less expensive. You still have to pay for CloudFront to transfer new versions of your files to edge locations, but you don't have to pay for invalidating files.

For more information about file versioning, see [Updating Existing Files Using Versioned File Names](#) (p. 105).

Determining Which Files to Invalidate

If you want to invalidate multiple files such as all of the files in a directory or all files that begin with the same characters, you can include the * wildcard at the end of the invalidation path. For more information about using the * wildcard, see [Invalidation paths](#).

If you want to invalidate selected files but your users don't necessarily access every file on your origin, you can determine which files viewers have requested from CloudFront and invalidate only those files. To determine which files viewers have requested, enable CloudFront access logging. For more information about access logs, see [Configuring and Using Standard Logs \(Access Logs\)](#) (p. 439).

Specifying the Files to Invalidate

Note the following about specifying the files that you want to invalidate.

Case sensitivity

Invalidation paths are case sensitive, so `/images/image.jpg` and `/images/Image.jpg` specify two different files.

Changing the URI Using a Lambda Function

If your CloudFront distribution triggers a Lambda function on viewer request events, and if the function changes the URI of the requested file, we recommend that you invalidate both URIs to remove the file from CloudFront edge caches:

- The URI in the viewer request
- The URI after the function changed it

For example, suppose your Lambda function changes the URI for a file from this:

```
https://d111111abcdef8.cloudfront.net/index.html
```

to a URI that includes a language directory:

```
https://d111111abcdef8.cloudfront.net/en/index.html
```

To invalidate the file, you must specify the following paths:

- `/index.html`
- `/en/index.html`

For more information, see [Invalidation paths](#).

Default root object

To invalidate the default root object (file), specify the path the same way that you specify the path for any other file.

Distribution types

You can invalidate only files that are associated with a web distribution. You cannot invalidate objects that are served by an RTMP distribution.

Forwarding cookies

If you configured CloudFront to forward cookies to your origin, CloudFront edge caches might contain several versions of the file. When you invalidate a file, CloudFront invalidates every cached version of the file regardless of its associated cookies. You can't selectively invalidate some versions and not others based on the associated cookies. For more information, see [Caching Content Based on Cookies \(p. 244\)](#).

Forwarding headers

If you configured CloudFront to forward a whitelist of headers to your origin and to cache based on the values of the headers, CloudFront edge caches might contain several versions of the file. When you invalidate a file, CloudFront invalidates every cached version of the file regardless of the header values. You can't selectively invalidate some versions and not others based on header values. (If you configure CloudFront to forward all headers to your origin, CloudFront doesn't cache your files.) For more information, see [Caching Content Based on Request Headers \(p. 246\)](#).

Forwarding query strings

If you configured CloudFront to forward query strings to your origin, you must include the query strings when invalidating files, as shown in the following examples:

- `/images/image.jpg?parameter1=a`
- `/images/image.jpg?parameter1=b`

If client requests include five different query strings for the same file, you can either invalidate the file five times, once for each query string, or you can use the `*` wildcard in the invalidation path, as shown in the following example:

```
/images/image.jpg*
```

For more information about using wildcards in the invalidation path, see [Invalidation paths](#). For more information about query strings, see [Caching Content Based on Query String Parameters \(p. 241\)](#). To determine which query strings are in use, you can enable CloudFront logging. For more information, see [Configuring and Using Standard Logs \(Access Logs\) \(p. 439\)](#).

Maximum Allowed

For information about the maximum number of invalidations allowed, see [Concurrent Invalidation Request Maximum \(p. 116\)](#).

Microsoft Smooth Streaming files

You cannot invalidate media files in the Microsoft Smooth Streaming format when you have enabled Smooth Streaming for the corresponding cache behavior.

Non-ASCII or unsafe characters in the path

If the path includes non-ASCII characters or unsafe characters as defined in [RFC 1783](#), URL-encode those characters. Do not URL-encode any other characters in the path, or CloudFront will not invalidate the old version of the updated file.

Invalidation paths

The path is relative to the distribution. For example, to invalidate the file at `https://d111111abcdef8.cloudfront.net/images/image2.jpg`, you would specify the following:

/images/image2.jpg

Note

In the [CloudFront console](#), you can omit the leading slash in the path, like this: images/image2.jpg. When you use the CloudFront API directly, invalidation paths must begin with a leading slash.

You can also invalidate multiple files simultaneously by using the * wildcard. The *, which replaces 0 or more characters, must be the last character in the invalidation path. Also, if you use the AWS command line interface (CLI) for invalidating files and you specify a path that includes the * wildcard, you must use quotes (") around the path (like this: "/*").

The following are some examples:

- To invalidate all of the files in a directory:

*/directory-path/**

- To invalidate a directory, all of its subdirectories, and all of the files in the directory and subdirectories:

*/directory-path**

- To invalidate all files that have the same name but different file name extensions, such as logo.jpg, logo.png, and logo.gif:

*/directory-path/file-name.**

- To invalidate all of the files in a directory for which the file name starts with the same characters (such as all of the files for a video in HLS format), regardless of the file name extension:

*/directory-path/initial-characters-in-file-name**

- When you configure CloudFront to cache based on query string parameters and you want to invalidate every version of a file:

*/directory-path/file-name.file-name-extension**

- To invalidate all of the files in a distribution:

*/**

The maximum length of a path is 4,000 characters. You can't use a wildcard within the path; only at the end of the path.

For information about invalidating files if you use a Lambda function to change the URI, see [Changing the URI Using a Lambda Function](#).

The charge to submit an invalidation path is the same regardless of the number of files you're invalidating: a single file (/images/logo.jpg) or all of the files that are associated with a distribution (/*). For more information, see [Amazon CloudFront Pricing](#).

If the invalidation path is a directory and if you have not standardized on a method for specifying directories—with or without a trailing slash (/)—we recommend that you invalidate the directory both with and without a trailing slash, for example, /images and /images/.

Signed URLs

If you are using signed URLs, invalidate a file by including only the portion of the URL before the question mark (?).

Invalidating Files Using the Console

You can use the CloudFront console to create and run an invalidation, display a list of the invalidations that you submitted previously, and display detailed information about an individual invalidation. You

can also copy an existing invalidation, edit the list of file paths, and run the edited invalidation. You can't remove invalidations from the list.

- [Invalidating Files \(p. 114\)](#)
- [Copying, Editing, and Rerunning an Existing Invalidation \(p. 114\)](#)
- [Canceling Invalidations \(p. 115\)](#)
- [Listing Invalidations \(p. 115\)](#)
- [Displaying Information about an Invalidation \(p. 115\)](#)

Invalidating Files

To invalidate files using the CloudFront console, do the following.

To invalidate files

1. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
2. Select the distribution for which you want to invalidate files.
3. Choose **Distribution Settings**.
4. Choose the **Invalidations** tab.
5. Choose **Create Invalidation**.
6. For the files that you want to invalidate, enter one invalidation path per line. For information about specifying invalidation paths, see [Specifying the Files to Invalidate \(p. 111\)](#).

Important

Specify file paths carefully. You can't cancel an invalidation request after you start it.

7. Choose **Invalidate**.

Copying, Editing, and Rerunning an Existing Invalidation

You can copy an invalidation that you created previously, update the list of invalidation paths, and run the updated invalidation. You cannot copy an existing invalidation, update the invalidation paths, and then save the updated invalidation without running it.

Important

If you copy an invalidation that is still in progress, update the list of invalidation paths, and then run the updated invalidation, CloudFront will not stop or delete the invalidation that you copied. If any invalidation paths appear in the original and in the copy, CloudFront will try to invalidate the files twice, and both invalidations will count against your maximum number of free invalidations for the month. If you've already reached the maximum number of free invalidations, you'll be charged for both invalidations of each file. For more information, see [Concurrent Invalidation Request Maximum \(p. 116\)](#).

To copy, edit, and rerun an existing invalidation

1. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
2. Select the distribution that contains the invalidation that you want to copy.
3. Choose **Distribution Settings**.
4. Choose the **Invalidations** tab.
5. Choose the invalidation that you want to copy.

If you aren't sure which invalidation you want to copy, you can choose an invalidation and choose **Details** to display detailed information about that invalidation.

6. Choose **Copy**.
7. Update the list of invalidation paths if applicable.
8. Choose **Invalidate**.

Canceling Invalidations

When you submit an invalidation request to CloudFront, CloudFront forwards the request to all edge locations within a few seconds, and each edge location starts processing the invalidation immediately. As a result, you can't cancel an invalidation after you submit it.

Listing Invalidations

You can display a list of the last 100 invalidations that you've created and run for a distribution by using the CloudFront console. If you want to get a list of more than 100 invalidations, use the `ListInvalidations` API action. For more information, see [ListInvalidations](#) in the *Amazon CloudFront API Reference*.

To list invalidations

1. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
2. Select the distribution for which you want to display a list of invalidations.
3. Choose **Distribution Settings**.
4. Choose the **Invalidations** tab.

Note

You can't remove invalidations from the list.

Displaying Information about an Invalidation

You can display detailed information about an invalidation, including distribution ID, invalidation ID, the status of the invalidation, the date and time that the invalidation was created, and a complete list of the invalidation paths.

To display information about an invalidation

1. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
2. Select the distribution that contains the invalidation that you want to display detailed information for.
3. Choose **Distribution Settings**.
4. Choose the **Invalidations** tab.
5. Choose the applicable invalidation.
6. Choose **Details**.

Invalidating Files Using the CloudFront API

For information about invalidating objects and about displaying information about invalidations using the CloudFront API, see the following topics in the *Amazon CloudFront API Reference*:

- Invalidating files: [CreateInvalidation](#)
- Getting a list of your invalidations: [ListInvalidations](#)
- Getting information about a specific invalidation: [GetInvalidation](#)

Concurrent Invalidation Request Maximum

If you're invalidating files individually, you can have invalidation requests for up to 3,000 files per distribution in progress at one time. This can be one invalidation request for up to 3,000 files, up to 3,000 requests for one file each, or any other combination that doesn't exceed 3,000 files. For example, you can submit 30 invalidation requests that invalidate 100 files each. As long as all 30 invalidation requests are still in progress, you can't submit any more invalidation requests. If you exceed the maximum, CloudFront returns an error message.

If you're using the * wildcard, you can have requests for up to 15 invalidation paths in progress at one time. You can also have invalidation requests for up to 3,000 individual files per distribution in progress at the same time; the maximum on wildcard invalidation requests allowed is independent of the maximum on invalidating files individually.

Paying for File Invalidation

The first 1,000 invalidation paths that you submit per month are free; you pay for each invalidation path over 1,000 in a month. An invalidation path can be for a single file (such as `/images/logo.jpg`) or for multiple files (such as `/images/*`). A path that includes the * wildcard counts as one path even if it causes CloudFront to invalidate thousands of files.

The maximum of 1,000 free invalidation paths per month applies to the total number of invalidation paths across all of the distributions that you create with one AWS account. For example, if you use the AWS account `john@example.com` to create three distributions, and you submit 600 invalidation paths for each distribution in a given month (for a total of 1,800 invalidation paths), AWS will charge you for 800 invalidation paths in that month. For specific information about invalidation pricing, see [Amazon CloudFront Pricing](#). For more information about invalidation paths, see [Invalidation paths](#).

Serving compressed files

You can use CloudFront to automatically compress files of certain types and serve the compressed files when viewers support them (viewers indicate their support for compressed files with the `Accept-Encoding` HTTP header). CloudFront can compress files using the Gzip and Brotli compression formats. When the viewer supports both formats, CloudFront uses Brotli.

Note

The Chrome and Firefox web browsers support Brotli compression only when the request is sent using HTTPS. These browsers do not support Brotli with HTTP requests.

CloudFront only compresses content using Brotli when both of the following are true:

- You configure the distribution to compress content. See the following section.
- You enable the **Cache Brotli objects** setting (set `EnableAcceptEncodingBrotli` to `true`). For more information, see [Cache compressed objects](#) (p. 86).

CloudFront compresses content using Gzip when just the first of the previous items is true.

When content is compressed, downloads can be faster because the files are smaller—in some cases, less than a quarter the size of the original. Especially for JavaScript and CSS files, faster downloads can result

in faster rendering of webpages for your users. In addition, because the cost of CloudFront data transfer is based on the total amount of data served, serving compressed files can be less expensive than serving uncompressed files.

If you're using a custom origin, you can configure your origin to compress files. Your origin might be able to compress files that CloudFront doesn't compress (see [File types that CloudFront compresses \(p. 119\)](#)). If your origin returns a compressed file to CloudFront, CloudFront detects that the file is compressed based on the value of the `Content-Encoding` header and doesn't compress the file again.

If you configure CloudFront to serve compressed content, you should also [cache compressed objects \(p. 86\)](#).

Configuring a CloudFront distribution to compress content

To configure a distribution to compress your content, update the cache behaviors that you want to serve the compressed content by using one of the following methods:

- **CloudFront console** – Update the **Compress objects automatically** setting. For more information, see [Creating a Distribution \(p. 37\)](#).
- **CloudFront API, AWS SDKs, AWS CLI, or AWS CloudFormation** – In a `CacheBehavior`, set the value of the `Compress` element to `true`. For more information, see [CreateDistribution](#) or [UpdateDistribution](#).

If you configure CloudFront to serve compressed content, you should also [cache compressed objects \(p. 86\)](#).

Using CloudFront to compress your content

CloudFront can compress files both for Amazon S3 origins and for custom origins. When you configure CloudFront to compress your content, you specify the setting in one or more cache behaviors.

When you configure CloudFront to compress your content, here's how CloudFront serves your content:

1. You create or update a CloudFront distribution and configure CloudFront to compress content. You also configure CloudFront to [cache compressed objects \(p. 86\)](#).
2. A viewer requests a file. The viewer includes the `Accept-Encoding` HTTP header in the request, and the header values include `gzip`, `br`, or both. This indicates that the viewer supports compressed content. When the viewer supports both formats, CloudFront uses Brotli.

Note

The Chrome and Firefox web browsers support Brotli compression only when the request is sent using HTTPS. These browsers do not support Brotli with HTTP requests.

CloudFront only compresses content using Brotli when both of the following are true:

- You configure the distribution to compress content. See the previous section.
- You enable the **Cache Brotli objects** setting (set `EnableAcceptEncodingBrotli` to `true`). For more information, see [Cache compressed objects \(p. 86\)](#).

CloudFront compresses content using Gzip when just the first of the previous items is true.

3. At the edge location, CloudFront checks the cache for a compressed version of the requested file.
4. If the compressed file is already in the cache, CloudFront returns the file to the viewer and skips the remaining steps.

5. If the compressed file is not in the cache, CloudFront sends the request to the origin server, which can be either an Amazon S3 bucket or a custom origin.

Note

If CloudFront has an uncompressed version of the file in the cache, it still sends a request to the origin.

6. The origin server returns an uncompressed version of the requested file to CloudFront.
7. CloudFront determines whether the file is compressible:
 - The file must be of a type that CloudFront compresses. For more information, see [File types that CloudFront compresses \(p. 119\)](#).
 - The file size must be between 1,000 and 10,000,000 bytes.
 - The response must include a `Content-Length` header so CloudFront can determine whether the size of the file is in the range that CloudFront compresses. If the `Content-Length` header is missing, CloudFront won't compress the file.
 - The response must not include a `Content-Encoding` header.
8. If the file is compressible, CloudFront compresses it, returns the compressed file to the viewer, and adds it to the cache.
9. The viewer uncompresses the file.

Note the following:

File types that CloudFront compresses

CloudFront compresses files for a large number of file types. For a complete list, see [File types that CloudFront compresses \(p. 119\)](#).

Size of files that CloudFront compresses

CloudFront compresses files that are between 1,000 bytes and 10,000,000 bytes in size.

Content-Length header

The origin must include a `Content-Length` header in the response so CloudFront can determine whether the size of the file is in the range that CloudFront compresses. If the `Content-Length` header is missing, CloudFront won't compress the file.

ETag header

When the uncompressed object from the origin includes a valid, strong `ETag` HTTP header, CloudFront converts the strong `ETag` header value to a weak `ETag`, and returns the weak `ETag` value to the viewer. Viewers can store the weak `ETag` value and use it to send conditional requests with the `If-None-Match` HTTP header. This allows viewers, CloudFront, and the origin to treat the compressed and uncompressed versions of an object as semantically equivalent, which reduces unnecessary data transfer.

A valid, strong `ETag` header value begins with a double quote character (`"`). To convert the strong `ETag` value to a weak one, CloudFront adds the characters `w/` to the beginning of the strong `ETag` value.

When the object from the origin includes a weak `ETag` header value (a value that begins with the characters `w/`), CloudFront does not modify this value, and returns it to the viewer as received from the origin.

When the object from the origin includes an invalid `ETag` header value (the value does not begin with `"` or with `w/`), CloudFront removes the `ETag` header and returns the object to the viewer without the `ETag` response header.

For more information, see the following pages in the MDN web docs:

- [Directives](#) (`ETag` HTTP header)

- [Weak validation](#) (HTTP conditional requests)
- [If-None-Match HTTP header](#)

Content already in edge locations when you configure CloudFront to compress files

CloudFront compresses files in each edge location when it gets the files from your origin. When you configure CloudFront to compress your content, CloudFront doesn't compress files that are already cached in edge locations. In addition, when a file expires in an edge location and CloudFront sends another request for the file to your origin, CloudFront doesn't compress the file if your origin returns an HTTP status code 304, which means that the edge location already has the latest version of the file. If you want CloudFront to compress the files that are already cached in edge locations, you need to invalidate those files. For more information, see [Invalidating Files \(p. 110\)](#).

Origin is already configured to compress files

If you configure CloudFront to compress files and the origin also compresses files, the origin includes a `Content-Encoding` header, which indicates that the file is already compressed. CloudFront returns the cached file to the viewer and caches it in the edge location.

Note

CloudFront does not compress a file if the response includes a `Content-Encoding` header, regardless of the header's value.

Request doesn't include the Accept-Encoding header

If the `Accept-Encoding` header is missing from the viewer request, CloudFront serves uncompressed content. If the `Accept-Encoding` header includes additional values such as `deflate` or `sdch`, CloudFront removes them before forwarding the request to the origin server.

Request uses HTTP 1.0

If a request to CloudFront uses HTTP 1.0, CloudFront removes the `Accept-Encoding` header and serves uncompressed content.

CloudFront is busy

In rare cases, when a CloudFront edge location is unusually busy, some files might not be compressed.

File types that CloudFront compresses

If you configure CloudFront to compress your content, CloudFront compresses files that have the following values in the `Content-Type` header:

- `application/dash+xml`
- `application/eot`
- `application/font`
- `application/font-sfnt`
- `application/javascript`
- `application/json`
- `application/opentype`
- `application/otf`
- `application/pkcs7-mime`
- `application/protobuf`
- `application/rss+xml`
- `application/truetype`
- `application/ttf`
- `application/vnd.apple.mpegurl`

- application/vnd.ms-fontobject
- application/xhtml+xml
- application/xml
- application/x-font-opentype
- application/x-font-truetype
- application/x-font-ttf
- application/x-httpd-cgi
- application/x-javascript
- application/x-mpegurl
- application/x-opentype
- application/x-otf
- application/x-perl
- application/x-ttf
- font/eot
- font/opentype
- font/otf
- font/ttf
- image/svg+xml
- text/css
- text/csv
- text/html
- text/javascript
- text/js
- text/plain
- text/richtext
- text/tab-separated-values
- text/xml
- text/x-component
- text/x-java-source
- text/x-script
- vnd.apple.mpegurl

Configuring Secure Access and Restricting Access to Content

For web distributions, CloudFront provides several options for securing content that it delivers. The following are some ways you can use CloudFront to secure and restrict access to content:

- Configure HTTPS connections
- Prevent users in specific geographic locations from accessing content
- Require users to access content using CloudFront signed URLs or signed cookies
- Set up field-level encryption for specific content fields
- Use AWS WAF to control access to your content

Topics

- [Using HTTPS with CloudFront \(p. 121\)](#)
- [Using Alternate Domain Names and HTTPS \(p. 132\)](#)
- [Serving Private Content with Signed URLs and Signed Cookies \(p. 145\)](#)
- [Restricting Access to Amazon S3 Content by Using an Origin Access Identity \(p. 209\)](#)
- [Using AWS WAF to Control Access to Your Content \(p. 216\)](#)
- [Restricting the Geographic Distribution of Your Content \(p. 217\)](#)
- [Using Field-Level Encryption to Help Protect Sensitive Data \(p. 219\)](#)

Using HTTPS with CloudFront

For web distributions, you can configure CloudFront to require that viewers use HTTPS to request your objects, so that connections are encrypted when CloudFront communicates with viewers. You also can configure CloudFront to use HTTPS to get objects from your origin, so that connections are encrypted when CloudFront communicates with your origin.

If you configure CloudFront to require HTTPS both to communicate with viewers and to communicate with your origin, here's what happens when CloudFront receives a request for an object:

1. A viewer submits an HTTPS request to CloudFront. There's some SSL/TLS negotiation here between the viewer and CloudFront. In the end, the viewer submits the request in an encrypted format.
2. If the object is in the CloudFront edge cache, CloudFront encrypts the response and returns it to the viewer, and the viewer decrypts it.
3. If the object is not in the CloudFront cache, CloudFront performs SSL/TLS negotiation with your origin and, when the negotiation is complete, forwards the request to your origin in an encrypted format.
4. Your origin decrypts the request, encrypts the requested object, and returns the object to CloudFront.
5. CloudFront decrypts the response, re-encrypts it, and forwards the object to the viewer. CloudFront also saves the object in the edge cache so that the object is available the next time it's requested.

6. The viewer decrypts the response.

The process works basically the same way whether your origin is an Amazon S3 bucket, MediaStore, or a custom origin such as an HTTP/S server.

Note

To help thwart SSL renegotiation-type attacks, CloudFront does not support renegotiation for viewer and origin requests.

For information about how to require HTTPS between viewers and CloudFront, and between CloudFront and your origin, see the following topics.

Topics

- [Requiring HTTPS for Communication Between Viewers and CloudFront \(p. 122\)](#)
- [Requiring HTTPS for Communication Between CloudFront and Your Custom Origin \(p. 123\)](#)
- [Requiring HTTPS for Communication Between CloudFront and Your Amazon S3 Origin \(p. 127\)](#)
- [Supported protocols and ciphers between viewers and CloudFront \(p. 128\)](#)
- [Supported protocols and ciphers between CloudFront and the origin \(p. 130\)](#)
- [Charges for HTTPS connections \(p. 132\)](#)

Requiring HTTPS for Communication Between Viewers and CloudFront

You can configure one or more cache behaviors in your CloudFront distribution to require HTTPS for communication between viewers and CloudFront. You also can configure one or more cache behaviors to allow both HTTP and HTTPS, so that CloudFront requires HTTPS for some objects but not for others. The configuration steps depend on which domain name you're using in object URLs:

- If you're using the domain name that CloudFront assigned to your distribution, such as `d111111abcdef8.cloudfront.net`, you change the **Viewer Protocol Policy** setting for one or more cache behaviors to require HTTPS communication. In that configuration, CloudFront provides the SSL/TLS certificate.

To change the value of **Viewer Protocol Policy** by using the CloudFront console, see the procedure later in this section.

For information about how to use the CloudFront API to change the value of the `ViewerProtocolPolicy` element, see [UpdateDistribution](#) in the *Amazon CloudFront API Reference*.

- If you're using your own domain name, such as `example.com`, you need to change several CloudFront settings. You also need to use an SSL/TLS certificate provided by AWS Certificate Manager (ACM), or import a certificate from a third-party certificate authority into ACM or the IAM certificate store. For more information, see [Using Alternate Domain Names and HTTPS \(p. 132\)](#).

Note

If you want to ensure that the objects that viewers get from CloudFront were encrypted when CloudFront got them from your origin, always use HTTPS between CloudFront and your origin. If you recently changed from HTTP to HTTPS between CloudFront and your origin, we recommend that you invalidate objects in CloudFront edge locations. CloudFront will return an object to a viewer regardless of whether the protocol used by the viewer (HTTP or HTTPS) matches the protocol that CloudFront used to get the object. For more information about removing or replacing objects in a distribution, see [Adding, Removing, or Replacing Content That CloudFront Distributes \(p. 104\)](#).

To require HTTPS between viewers and CloudFront for one or more cache behaviors, perform the following procedure.

To configure CloudFront to require HTTPS between viewers and CloudFront

1. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
2. In the top pane of the CloudFront console, choose the ID for the distribution that you want to update.
3. On the **Behaviors** tab, choose the cache behavior that you want to update, and then choose **Edit**.
4. Specify one of the following values for **Viewer Protocol Policy**:

Redirect HTTP to HTTPS

Viewers can use both protocols. HTTP `GET` and `HEAD` requests are automatically redirected to HTTPS requests. CloudFront returns HTTP status code 301 (Moved Permanently) along with the new HTTPS URL. The viewer then resubmits the request to CloudFront using the HTTPS URL.

Important

If you send `POST`, `PUT`, `DELETE`, `OPTIONS`, or `PATCH` over HTTP with an HTTP to HTTPS cache behavior and a request protocol version of HTTP 1.1 or above, CloudFront redirects the request to a HTTPS location with a HTTP status code 307 (Temporary Redirect). This guarantees that the request is sent again to the new location using the same method and body payload.

If you send `POST`, `PUT`, `DELETE`, `OPTIONS`, or `PATCH` requests over HTTP to HTTPS cache behavior with a request protocol version below HTTP 1.1, CloudFront returns a HTTP status code 403 (Forbidden).

When a viewer makes an HTTP request that is redirected to an HTTPS request, CloudFront charges for both requests. For the HTTP request, the charge is only for the request and for the headers that CloudFront returns to the viewer. For the HTTPS request, the charge is for the request, and for the headers and the object that are returned by your origin.

HTTPS Only

Viewers can access your content only if they're using HTTPS. If a viewer sends an HTTP request instead of an HTTPS request, CloudFront returns HTTP status code 403 (Forbidden) and does not return the object.

5. Choose **Yes, Edit**.
6. Repeat steps 3 through 5 for each additional cache behavior that you want to require HTTPS for between viewers and CloudFront.
7. Confirm the following before you use the updated configuration in a production environment:
 - The path pattern in each cache behavior applies only to the requests that you want viewers to use HTTPS for.
 - The cache behaviors are listed in the order that you want CloudFront to evaluate them in. For more information, see [Path Pattern \(p. 47\)](#).
 - The cache behaviors are routing requests to the correct origins.

Requiring HTTPS for Communication Between CloudFront and Your Custom Origin

If you want to require HTTPS for communication between CloudFront and your custom origin, the steps you take depend on whether you're using the domain name that CloudFront assigned to

your distribution (like `d111111abcdef8.cloudfront.net`) or your own alternate domain name (like `example.com`).

Note

If your custom origin is an Amazon S3 bucket that's configured as a website endpoint, you can't configure CloudFront to use HTTPS with your origin because Amazon S3 doesn't support HTTPS for website endpoints.

Use the default CloudFront domain name

If you're using the domain name that CloudFront assigned to your distribution in the URLs for your objects (for example, `https://d111111abcdef8.cloudfront.net/logo.jpg`), you can require HTTPS by following the procedures in this topic to do the following:

- Change the **Origin Protocol Policy** setting for specific origins in your distribution
- Install an SSL/TLS certificate on your custom origin server (this isn't required when you use an Amazon S3 origin)

Use an alternate domain name

Instead of using the default domain name with your distribution, you can add an alternate domain name that's easier to work with, like `example.com`.

To require HTTPS for communication when you use an alternate domain name, follow the steps and guidance in [Using Alternate Domain Names and HTTPS \(p. 132\)](#).

Topics

- [Changing CloudFront Settings \(p. 124\)](#)
- [Installing an SSL/TLS Certificate on Your Custom Origin Server \(p. 125\)](#)
- [About RSA and ECDSA Ciphers \(p. 126\)](#)

Changing CloudFront Settings

The following procedure explains how to configure CloudFront to use HTTPS to communicate with an Elastic Load Balancing load balancer, an Amazon EC2 instance, or another custom origin. For information about using the CloudFront API to update a web distribution, see [UpdateDistribution](#) in the *Amazon CloudFront API Reference*.

To configure CloudFront to require HTTPS between CloudFront and your custom origin

1. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
2. In the top pane of the CloudFront console, choose the ID for the distribution that you want to update.
3. On the **Origins** tab, choose the origin that you want to update, and then choose **Edit**.
4. Update the following settings:

Origin Protocol Policy

Change the **Origin Protocol Policy** for the applicable origins in your distribution:

- **HTTPS Only** – CloudFront uses only HTTPS to communicate with your custom origin.
- **Match Viewer** – CloudFront communicates with your custom origin using HTTP or HTTPS, depending on the protocol of the viewer request. For example, if you choose **Match Viewer** for **Origin Protocol Policy** and the viewer uses HTTPS to request an object from CloudFront, CloudFront also uses HTTPS to forward the request to your origin.

Choose **Match Viewer** only if you specify **Redirect HTTP to HTTPS** or **HTTPS Only** for **Viewer Protocol Policy**.

CloudFront caches the object only once even if viewers make requests using both HTTP and HTTPS protocols.

Origin SSL Protocols

Choose the **Origin SSL Protocols** for the applicable origins in your distribution. The SSLv3 protocol is less secure, so we recommend that you choose SSLv3 only if your origin doesn't support TLSv1 or later. The TLSv1 handshake is both backwards and forwards compatible with SSLv3, but TLSv1.1 and TLSv1.2 are not. When you choose SSLv3, CloudFront *only* sends SSLv3 handshake requests.

5. Choose **Yes, Edit**.
6. Repeat steps 3 through 5 for each additional origin that you want to require HTTPS for between CloudFront and your custom origin.
7. Confirm the following before you use the updated configuration in a production environment:
 - The path pattern in each cache behavior applies only to the requests that you want viewers to use HTTPS for.
 - The cache behaviors are listed in the order that you want CloudFront to evaluate them in. For more information, see [Path Pattern](#) (p. 47).
 - The cache behaviors are routing requests to the origins that you changed the **Origin Protocol Policy** for.

Installing an SSL/TLS Certificate on Your Custom Origin Server

You can use an SSL/TLS certificate from the following sources on your custom origin:

- If your origin is an Elastic Load Balancing load balancer, you can use a certificate provided by AWS Certificate Manager (ACM). You also can use a certificate that is signed by a trusted third-party certificate authority and imported into ACM.
- For origins other than Elastic Load Balancing load balancers, you must use a certificate that is signed by a trusted third-party certificate authority (CA), for example, Comodo, DigiCert, or Symantec.

When CloudFront uses HTTPS to communicate with your origin, CloudFront verifies that the certificate was issued by a trusted certificate authority. CloudFront supports the same certificate authorities that Mozilla does. For the current list, see [Mozilla Included CA Certificate List](#). You can't use a self-signed certificate for HTTPS communication between CloudFront and your origin.

Important

If the origin server returns an expired certificate, an invalid certificate, or a self-signed certificate, or if the origin server returns the certificate chain in the wrong order, CloudFront drops the TCP connection, returns HTTP status code 502 (Bad Gateway), and sets the `x-cache` header to `Error from cloudfront`. Also, if the full chain of certificates, including the intermediate certificate, is not present, CloudFront drops the TCP connection.

The certificate returned from the origin must cover the domain that you specified for **Origin Domain Name** for the corresponding origin in your distribution. In addition, if you configured CloudFront to forward the `Host` header to your origin, the origin must respond with a certificate matching the domain in the `Host` header.

About RSA and ECDSA Ciphers

The encryption strength of a communications connection depends on the key size and strength of the algorithm that you choose for your origin server's certificate. The two options that CloudFront supports for connections with a custom origin are RSA and Elliptic Curve Digital Signature Algorithm (ECDSA).

For lists of the RSA and ECDSA ciphers supported by CloudFront, see [Supported protocols and ciphers between CloudFront and the origin](#) (p. 130).

How RSA Ciphers Work

CloudFront and origin servers typically use RSA 2048-bit asymmetric keys for SSL/TLS termination. RSA algorithms use the product of two large prime numbers, with another number added to it to create a public key. The private key is a related number. The strength of RSA relies on the presumed difficulty of breaking a key that requires factoring the product of two large prime numbers. However, improvements in computer technology have weakened RSA algorithms because faster computer calculations mean that it's now easier to break the encryption.

If you want to maintain encryption strength while continuing to use RSA, one option would be to increase the size of your RSA keys. However, this approach isn't easily scalable because using larger keys increases the compute cost for cryptography.

How ECDSA Ciphers Work

Alternatively, you could use an ECDSA certificate. ECDSA bases its security on a more complex mathematical problem than RSA that is harder to solve, which means that it takes more computer processing time to break ECDSA encryption. ECDSA is built on the principle that it is difficult to solve for the discrete logarithm of a random elliptic curve when its base is known, also known as the Elliptic Curve Discrete Logarithm Problem (ECDLP). This means that you can use shorter key lengths to achieve the equivalent security of using RSA with much larger key sizes.

In addition to providing better security, using ECDSA's smaller keys enables faster computing of algorithms, smaller digital certificates, and fewer bits to transmit during the SSL/TLS handshake. As a result, the smaller keys reduce the time that it takes for you to create and sign digital certificates for SSL/TLS termination on origin servers. Using a smaller key size therefore can increase throughput by reducing the compute cycles needed for cryptography, freeing up server resources to process other work.

Choosing Between RSA and ECDSA Ciphers

Sample tests that we have run to compare, for example, 2048-bit RSA to 256-bit ECDSA (nistp256) have indicated that the nistp256 option was 95% faster than 2048-bit RSA while providing the same security strength as 3072-bit RSA.

CloudFront continues to support RSA for SSL/TLS connections. However, if you have concerns about the strength of your current encryption for SSL/TLS authentication for your origin servers, ECDSA could be a better option. The effort to enable ECDSA digital certificates compared to the security benefit that ECDSA brings is a trade-off that you will have to weigh in making your decision. In addition to enabling stronger encryption, the reduction in computational cost of cryptography while using ECDSA at your origin servers is an added advantage.

Using ECDSA Ciphers

To use ECDSA for communications between CloudFront and your origin, do the following:

1. Generate a private key by using one of the supported curves (prime256v1, secp384r1, or X25519).
2. Generate an ECDSA Digital Certificate in the X.509 PEM format with a trusted certificate authority.
3. Set up your origin to prefer the ECDSA certificate.

Using ECDSA doesn't require any settings changes in the CloudFront console or APIs, and there is no additional fee.

Requiring HTTPS for Communication Between CloudFront and Your Amazon S3 Origin

When your origin is an Amazon S3 bucket, your options for using HTTPS for communications with CloudFront depend on how you're using the bucket. If your Amazon S3 bucket is configured as a website endpoint, you can't configure CloudFront to use HTTPS to communicate with your origin because Amazon S3 doesn't support HTTPS connections in that configuration.

When your origin is an Amazon S3 bucket that supports HTTPS communication, CloudFront always forwards requests to S3 by using the protocol that viewers used to submit the requests. The default setting for the [Origin Protocol Policy](#) (p. 45) setting is **Match Viewer** and can't be changed.

If you want to require HTTPS for communication between CloudFront and Amazon S3, you must change the value of **Viewer Protocol Policy** to **Redirect HTTP to HTTPS** or **HTTPS Only**. The procedure later in this section explains how to use the CloudFront console to change **Viewer Protocol Policy**. For information about using the CloudFront API to update the `ViewerProtocolPolicy` element for a web distribution, see [UpdateDistribution](#) in the *Amazon CloudFront API Reference*.

When you use HTTPS with an Amazon S3 bucket that supports HTTPS communication, Amazon S3 provides the SSL/TLS certificate, so you don't have to.

To configure CloudFront to require HTTPS to your Amazon S3 origin

1. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
2. In the top pane of the CloudFront console, choose the ID for the distribution that you want to update.
3. On the **Behaviors** tab, choose the cache behavior that you want to update, and then choose **Edit**.
4. Specify one of the following values for **Viewer Protocol Policy**:

Redirect HTTP to HTTPS

Viewers can use both protocols, but HTTP requests are automatically redirected to HTTPS requests. CloudFront returns HTTP status code 301 (Moved Permanently) along with the new HTTPS URL. The viewer then resubmits the request to CloudFront using the HTTPS URL.

Important

CloudFront doesn't redirect `DELETE`, `OPTIONS`, `PATCH`, `POST`, or `PUT` requests from HTTP to HTTPS. If you configure a cache behavior to redirect to HTTPS, CloudFront responds to HTTP `DELETE`, `OPTIONS`, `PATCH`, `POST`, or `PUT` requests for that cache behavior with HTTP status code 403 (Forbidden).

When a viewer makes an HTTP request that is redirected to an HTTPS request, CloudFront charges for both requests. For the HTTP request, the charge is only for the request and for the headers that CloudFront returns to the viewer. For the HTTPS request, the charge is for the request, and for the headers and the object returned by your origin.

HTTPS Only

Viewers can access your content only if they're using HTTPS. If a viewer sends an HTTP request instead of an HTTPS request, CloudFront returns HTTP status code 403 (Forbidden) and does not return the object.

5. Choose **Yes, Edit**.

6. Repeat steps 3 through 5 for each additional cache behavior that you want to require HTTPS for between viewers and CloudFront, and between CloudFront and S3.
7. Confirm the following before you use the updated configuration in a production environment:
 - The path pattern in each cache behavior applies only to the requests that you want viewers to use HTTPS for.
 - The cache behaviors are listed in the order that you want CloudFront to evaluate them in. For more information, see [Path Pattern \(p. 47\)](#).
 - The cache behaviors are routing requests to the correct origins.

Supported protocols and ciphers between viewers and CloudFront

When you [require HTTPS between viewers and your CloudFront distribution \(p. 49\)](#), you must choose a [security policy \(p. 56\)](#), which determines the following settings.

- The minimum SSL/TLS protocol that CloudFront uses to communicate with viewers.
- The ciphers that CloudFront can use to encrypt the communication with viewers.

To choose a security policy, specify the applicable value for [Security Policy \(p. 56\)](#). The following table lists the protocols and ciphers that CloudFront can use for each security policy.

A viewer must support at least one of the supported ciphers to establish an HTTPS connection with CloudFront. If you're using an SSL/TLS certificate in AWS Certificate Manager, a viewer must support one of the *-RSA-* ciphers. CloudFront chooses a cipher in the listed order from among the ciphers that the viewer supports. See also [OpenSSL, s2n, and RFC cipher names \(p. 129\)](#).

	Security policy					
SSL/TLS protocols supported						
TLSv1.3 ¹	◆	◆	◆	◆	◆	◆
TLSv1.2	◆	◆	◆	◆	◆	◆
TLSv1.1	◆	◆	◆	◆		
TLSv1	◆	◆	◆			
SSLv3	◆					
Ciphers supported						
TLS_AES_128_GCM_SHA256	◆	◆	◆	◆	◆	◆
TLS_AES_256_GCM_SHA384	◆	◆	◆	◆	◆	◆
TLS_CHACHA20_POLY1305_SHA256		◆	◆	◆	◆	◆
ECDHE-RSA-AES128-GCM-SHA256	◆	◆	◆	◆	◆	◆
ECDHE-RSA-AES128-SHA256	◆	◆	◆	◆	◆	◆
ECDHE-RSA-AES128-SHA	◆	◆	◆	◆		

	Security policy					
ECDHE-RSA-AES256-GCM-SHA384	◆	◆	◆	◆	◆	◆
ECDHE-RSA-CHACHA20-POLY1305	◆	◆	◆	◆	◆	◆
ECDHE-RSA-AES256-SHA384	◆	◆	◆	◆	◆	◆
ECDHE-RSA-AES256-SHA	◆	◆	◆	◆		
AES128-GCM-SHA256	◆	◆	◆	◆	◆	
AES256-GCM-SHA384	◆	◆	◆	◆	◆	
AES128-SHA256	◆	◆	◆	◆	◆	
AES256-SHA	◆	◆	◆	◆		
AES128-SHA	◆	◆	◆	◆		
DES-CBC3-SHA	◆	◆				
RC4-MD5	◆					

¹CloudFront supports one round trip time (1-RTT) handshakes for TLSv1.3, but does not support zero round trip time (0-RTT) handshakes.

OpenSSL, s2n, and RFC cipher names

OpenSSL and s2n use different names for ciphers than the TLS standards use ([RFC 2246](#), [RFC 4346](#), [RFC 5246](#), and [RFC 8446](#)). The following table maps the OpenSSL and s2n names to the RFC name for each cipher.

For all elliptic curve ciphers, CloudFront supports the following elliptic curves:

- prime256v1
- secp384r1
- X25519

OpenSSL and s2n cipher name	RFC cipher name
TLS_AES_128_GCM_SHA256	TLS_AES_128_GCM_SHA256
TLS_AES_256_GCM_SHA384	TLS_AES_256_GCM_SHA384
TLS_CHACHA20_POLY1305_SHA256	TLS_CHACHA20_POLY1305_SHA256
ECDHE-RSA-AES128-GCM-SHA256	TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
ECDHE-RSA-AES128-SHA256	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256
ECDHE-RSA-AES128-SHA	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA
ECDHE-RSA-AES256-GCM-SHA384	TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
ECDHE-RSA-CHACHA20-POLY1305	TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256

OpenSSL and s2n cipher name	RFC cipher name
ECDHE-RSA-AES256-SHA384	TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384
ECDHE-RSA-AES256-SHA	TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA
AES128-GCM-SHA256	TLS_RSA_WITH_AES_128_GCM_SHA256
AES256-GCM-SHA384	TLS_RSA_WITH_AES_256_GCM_SHA384
AES128-SHA256	TLS_RSA_WITH_AES_128_CBC_SHA256
AES256-SHA	TLS_RSA_WITH_AES_256_CBC_SHA
AES128-SHA	TLS_RSA_WITH_AES_128_CBC_SHA
DES-CBC3-SHA	TLS_RSA_WITH_3DES_EDE_CBC_SHA
RC4-MD5	TLS_RSA_WITH_RC4_128_MD5

Supported signature schemes between viewers and CloudFront

CloudFront supports the following signature schemes for connections between viewers and CloudFront.

- TLS_SIGNATURE_SCHEME_RSA_PSS_RSAE_SHA256
- TLS_SIGNATURE_SCHEME_RSA_PSS_RSAE_SHA384
- TLS_SIGNATURE_SCHEME_RSA_PSS_RSAE_SHA512
- TLS_SIGNATURE_SCHEME_RSA_PKCS1_SHA256
- TLS_SIGNATURE_SCHEME_RSA_PKCS1_SHA384
- TLS_SIGNATURE_SCHEME_RSA_PKCS1_SHA512
- TLS_SIGNATURE_SCHEME_RSA_PKCS1_SHA224
- TLS_SIGNATURE_SCHEME_RSA_PKCS1_SHA1

Supported protocols and ciphers between CloudFront and the origin

If you choose to [require HTTPS between CloudFront and your origin](#), you can decide [which SSL/TLS protocol to allow](#) for the secure connection, and then pick any supported cipher for CloudFront (see the following tables) to establish an HTTPS connection to your origin.

CloudFront can forward HTTPS requests to the origin server by using the ECDSA or RSA ciphers listed in the following tables. Your origin server must support at least one of these ciphers for CloudFront to establish an HTTPS connection to your origin. To learn more about the two types of ciphers that CloudFront supports, see [About RSA and ECDSA Ciphers \(p. 126\)](#).

OpenSSL and s2n use different names for ciphers than the TLS standards use ([RFC 2246](#), [RFC 4346](#), [RFC 5246](#), and [RFC 8446](#)). The following tables map the OpenSSL and s2n names to the RFC name for each cipher.

Supported RSA ciphers

CloudFront supports the following RSA ciphers for connections with an origin.

For all elliptic curve ciphers, CloudFront supports the following elliptic curves:

- prime256v1
- secp384r1
- X25519

OpenSSL and s2n cipher name	RFC cipher name
ECDHE-RSA-AES256-GCM-SHA384	TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
ECDHE-RSA-AES256-SHA384	TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384
ECDHE-RSA-AES256-SHA	TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA
ECDHE-RSA-AES128-GCM-SHA256	TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
ECDHE-RSA-AES128-SHA256	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256
ECDHE-RSA-AES128-SHA	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA
AES256-SHA	TLS_RSA_WITH_AES_256_CBC_SHA
AES128-SHA	TLS_RSA_WITH_AES_128_CBC_SHA
DES-CBC3-SHA	TLS_RSA_WITH_3DES_EDE_CBC_SHA
RC4-MD5	TLS_RSA_WITH_RC4_128_MD5

Supported ECDSA ciphers

CloudFront supports the following ECDSA ciphers for connections with an origin.

For all elliptic curve ciphers, CloudFront supports the following elliptic curves:

- prime256v1
- secp384r1
- X25519

OpenSSL and s2n cipher name	RFC cipher name
ECDHE-ECDSA-AES256-GCM-SHA384	TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384
ECDHE-ECDSA-AES256-SHA384	TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384
ECDHE-ECDSA-AES256-SHA	TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA
ECDHE-ECDSA-AES128-GCM-SHA256	TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
ECDHE-ECDSA-AES128-SHA256	TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256
ECDHE-ECDSA-AES128-SHA	TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA

Supported signature schemes between CloudFront and the origin

CloudFront supports the following signature schemes for connections between CloudFront and the origin.

- TLS_SIGNATURE_SCHEME_RSA_PKCS1_SHA256

- TLS_SIGNATURE_SCHEME_RSA_PKCS1_SHA384
- TLS_SIGNATURE_SCHEME_RSA_PKCS1_SHA512
- TLS_SIGNATURE_SCHEME_RSA_PKCS1_SHA224
- TLS_SIGNATURE_SCHEME_ECDSA_SHA256
- TLS_SIGNATURE_SCHEME_ECDSA_SHA384
- TLS_SIGNATURE_SCHEME_ECDSA_SHA512
- TLS_SIGNATURE_SCHEME_ECDSA_SHA224
- TLS_SIGNATURE_SCHEME_RSA_PKCS1_SHA1
- TLS_SIGNATURE_SCHEME_ECDSA_SHA1

Charges for HTTPS connections

You always incur a surcharge for HTTPS requests. For more information, see [Amazon CloudFront Pricing](#).

Using Alternate Domain Names and HTTPS

If you want to use your own domain name in the URLs for your files (for example, `https://www.example.com/image.jpg`) and you want your viewers to use HTTPS, you must complete the steps in this topic. (If you use the default CloudFront distribution domain name in your URLs, for example, `https://d111111abcdef8.cloudfront.net/image.jpg`, follow the guidance in the following topic instead: [Requiring HTTPS for Communication Between Viewers and CloudFront \(p. 122\)](#).)

Important

When you add a certificate to your distribution, CloudFront immediately propagates the certificate to all of its edge locations. As new edge locations become available, CloudFront propagates the certificate to those locations, too. You can't restrict the edge locations that CloudFront propagates the certificates to.

Topics

- [Choosing How CloudFront Serves HTTPS Requests \(p. 132\)](#)
- [Requirements for Using SSL/TLS Certificates with CloudFront \(p. 134\)](#)
- [Quotas on Using SSL/TLS Certificates with CloudFront \(HTTPS Between Viewers and CloudFront Only\) \(p. 137\)](#)
- [Configuring Alternate Domain Names and HTTPS \(p. 139\)](#)
- [Determining the Size of the Public Key in an SSL/TLS Certificate \(p. 142\)](#)
- [Increasing the Quotas for SSL/TLS Certificates \(p. 142\)](#)
- [Rotating SSL/TLS Certificates \(p. 143\)](#)
- [Reverting from a Custom SSL/TLS Certificate to the Default CloudFront Certificate \(p. 144\)](#)
- [Switching from a Custom SSL/TLS Certificate with Dedicated IP Addresses to SNI \(p. 145\)](#)

Choosing How CloudFront Serves HTTPS Requests

If you want your viewers to use HTTPS and to use alternate domain names for your files, choose one of the following options for how CloudFront serves HTTPS requests:

- Use Server Name Indication (SNI) – Recommended
- Use a dedicated IP address in each edge location

This section explains how each option works.

Using SNI to Serve HTTPS Requests (Works for Most Clients)

Server Name Indication (SNI) is an extension to the TLS protocol that is supported by browsers and clients released after 2010. If you configure CloudFront to serve HTTPS requests using SNI, CloudFront associates your alternate domain name with an IP address for each edge location. When a viewer submits an HTTPS request for your content, DNS routes the request to the IP address for the correct edge location. The IP address to your domain name is determined during the SSL/TLS handshake negotiation; the IP address isn't dedicated to your distribution.

The SSL/TLS negotiation occurs early in the process of establishing an HTTPS connection. If CloudFront can't immediately determine which domain the request is for, it drops the connection. When a viewer that supports SNI submits an HTTPS request for your content, here's what happens:

1. The viewer automatically gets the domain name from the request URL and adds it to a field in the request header.
2. When CloudFront receives the request, it finds the domain name in the request header and responds to the request with the applicable SSL/TLS certificate.
3. The viewer and CloudFront perform SSL/TLS negotiation.
4. CloudFront returns the requested content to the viewer.

For a current list of the browsers that support SNI, see the Wikipedia entry [Server Name Indication](#).

If you want to use SNI but some of your users' browsers don't support SNI, you have several options:

- Configure CloudFront to serve HTTPS requests by using dedicated IP addresses instead of SNI. For more information, see [Using a Dedicated IP Addresses to Serve HTTPS Requests \(Works for All Clients\)](#) (p. 133).
- Use the CloudFront SSL/TLS certificate instead of a custom certificate. This requires that you use the CloudFront domain name for your distribution in the URLs for your files, for example, `https://d111111abcdef8.cloudfront.net/logo.png`.

If you use the default CloudFront certificate, viewers must support the SSL protocol TLSv1 or later. CloudFront doesn't support SSLv3 with the default CloudFront certificate.

You also must change the SSL/TLS certificate that CloudFront is using from a custom certificate to the default CloudFront certificate:

- If you haven't used your distribution to distribute your content, you can just change the configuration. For more information, see [Updating a Distribution](#) (p. 63).
- If you have used your distribution to distribute your content, you must create a new CloudFront distribution and change the URLs for your files to reduce or eliminate the amount of time that your content is unavailable. For more information, see [Reverting from a Custom SSL/TLS Certificate to the Default CloudFront Certificate](#) (p. 144).
- If you can control which browser your users use, have them upgrade their browser to one that supports SNI.
- Use HTTP instead of HTTPS.

Using a Dedicated IP Addresses to Serve HTTPS Requests (Works for All Clients)

Server Name Indication (SNI) is one way to associate a request with a domain. Another way is to use a dedicated IP address. If you have users who can't upgrade to a browser or client released after 2010, you can use a dedicated IP address to serve HTTPS requests. For a current list of the browsers that support SNI, see the Wikipedia entry [Server Name Indication](#).

Important

If you configure CloudFront to serve HTTPS requests using dedicated IP addresses, you incur an additional monthly charge. The charge begins when you associate your SSL/TLS certificate with a distribution and you enable the distribution. For more information about CloudFront pricing, see [Amazon CloudFront Pricing](#). In addition, see [Using the Same Certificate for Multiple CloudFront Distributions](#) (p. 138).

When you configure CloudFront to serve HTTPS requests using dedicated IP addresses, CloudFront associates your alternate domain name with a dedicated IP address in each CloudFront edge location. When a viewer submits an HTTPS request for your content, here's what happens:

1. DNS routes the request to the IP address for your distribution in the applicable edge location.
2. CloudFront uses the IP address to identify your distribution and to determine which SSL/TLS certificate to return to the viewer.
3. The viewer and CloudFront perform SSL/TLS negotiation using your SSL/TLS certificate.
4. CloudFront returns the requested content to the viewer.

This method works for every HTTPS request, regardless of the browser or other viewer that the user is using.

Requirements for Using SSL/TLS Certificates with CloudFront

The requirements for SSL/TLS certificates are described in this topic. They apply, except as noted, to both of the following:

- Certificates for using HTTPS between viewers and CloudFront
- Certificates for using HTTPS between CloudFront and your origin

Topics

- [Certificate Issuer](#) (p. 134)
- [AWS Region that You Request a Certificate In \(for AWS Certificate Manager\)](#) (p. 135)
- [Certificate Format](#) (p. 135)
- [Intermediate Certificates](#) (p. 135)
- [Key Type](#) (p. 136)
- [Private Key](#) (p. 136)
- [Permissions](#) (p. 136)
- [Size of the Public Key](#) (p. 136)
- [Supported Types of Certificates](#) (p. 136)
- [Certificate Expiration Date and Renewal](#) (p. 137)
- [Domain Names in the CloudFront Distribution and in the Certificate](#) (p. 137)
- [Minimum SSL Protocol Version](#) (p. 137)
- [Supported HTTP Versions](#) (p. 137)

Certificate Issuer

The certificate issuer you must use depends on whether you want to require HTTPS between viewers and CloudFront or between CloudFront and your origin:

- **HTTPS between viewers and CloudFront** – You can use a certificate that was issued by a trusted certificate authority (CA) such as Comodo, DigiCert, or Symantec, or you can use a certificate provided by AWS Certificate Manager (ACM).

Important

If you want to use an alternate domain name with your CloudFront distribution, you must verify to CloudFront that you have authorized rights to use the alternate domain name. To do this, you must attach a valid certificate to your distribution, and make sure that the certificate comes from a trusted CA that is listed on the [Mozilla Included CA Certificate List](#). CloudFront does not allow you to use a self-signed certificate to verify your authorized rights to use an alternate domain name.

- **HTTPS between CloudFront and a custom origin** – If the origin is *not* an Elastic Load Balancing (ELB) load balancer, such as Amazon EC2, the certificate must be issued by a trusted CA such as Comodo, DigiCert, or Symantec. If your origin is an ELB load balancer, you can also use a certificate provided by ACM.

Important

When CloudFront uses HTTPS to communicate with your origin, CloudFront verifies that the certificate was issued by a trusted CA. CloudFront supports the same certificate authorities as Mozilla; for the current list, see [Mozilla Included CA Certificate List](#). You cannot use a self-signed certificate for HTTPS communication between CloudFront and your origin.

For more information about getting and installing an SSL/TLS certificate, refer to the documentation for your HTTP server software and to the documentation for the certificate authority. For information about ACM, see the [AWS Certificate Manager User Guide](#).

AWS Region that You Request a Certificate In (for AWS Certificate Manager)

If you want to require HTTPS between viewers and CloudFront, you must change the AWS Region to US East (N. Virginia) in the AWS Certificate Manager console before you request or import a certificate.

If you want to require HTTPS between CloudFront and your origin, and you're using an ELB load balancer as your origin, you can request or import a certificate in any Region.

Certificate Format

The certificate must be in X.509 PEM format. This is the default format if you're using AWS Certificate Manager.

Intermediate Certificates

If you're using a third-party certificate authority (CA), in the .pem file, list all of the intermediate certificates in the certificate chain, beginning with one for the CA that signed the certificate for your domain. Typically, you'll find a file on the CA website that lists intermediate and root certificates in the proper chained order.

Important

Do not include the following: the root certificate, intermediate certificates that are not in the trust path, or your CA's public key certificate.

Here's an example:

```
-----BEGIN CERTIFICATE-----  
Intermediate certificate 2  
-----END CERTIFICATE-----  
-----BEGIN CERTIFICATE-----
```



```
Intermediate certificate 1  
-----END CERTIFICATE-----
```

Key Type

CloudFront supports only RSA public/private key pairs.

Private Key

If you're using a certificate from a third-party certificate authority (CA), note the following:

- The private key must match the public key that is in the certificate.
- The private key also must be an RSA private key in PEM format, where the PEM header is `BEGIN RSA PRIVATE KEY` and the footer is `END RSA PRIVATE KEY`.
- The private key cannot be encrypted with a password.

If AWS Certificate Manager (ACM) provided the certificate, ACM doesn't release the private key. The private key is stored in ACM for use by AWS services that are integrated with ACM.

Permissions

You must have permission to use and import the SSL/TLS certificate. If you're using AWS Certificate Manager (ACM), we recommend that you use AWS Identity and Access Management permissions to restrict access to the certificates. For more information, see [Permissions and Policies](#) in the *AWS Certificate Manager User Guide*.

Size of the Public Key

The length of the public key for a certificate depends on where you're storing it.

- Importing a certificate into AWS Certificate Manager (ACM): public key length must be 1024 bits or 2048 bits. The maximum length for a certificate that you use with CloudFront is 2048 bits, even though ACM supports larger keys.
- Uploading a certificate to the AWS Identity and Access Management (IAM) certificate store: maximum size of the public key is 2048 bits.

We recommend using 2048 bits.

For information about the public keys for certificates provided by ACM, see [ACM Certificate Characteristics](#) in the *AWS Certificate Manager User Guide*.

For information about how to determine the size of the public key, see [Determining the Size of the Public Key in an SSL/TLS Certificate \(p. 142\)](#).

Supported Types of Certificates

CloudFront supports all types of certificates, including the following:

- Domain-validated certificates
- Extended validation (EV) certificates
- High-assurance certificates
- Wildcard certificates (`*.example.com`)
- Subject alternative name (SAN) certificates (`example.com` and `example.net`)

Certificate Expiration Date and Renewal

If you're using certificates that you get from a third-party certificate authority (CA), you must monitor certificate expiration dates and renew SSL/TLS certificates that you import into AWS Certificate Manager (ACM) or upload to the AWS Identity and Access Management certificate store.

If you're using ACM-provided certificates, ACM manages certificate renewals for you. For more information, see [Managed Renewal](#) in the *AWS Certificate Manager User Guide*.

Domain Names in the CloudFront Distribution and in the Certificate

When you're using a custom origin, the SSL/TLS certificate on your origin includes a domain name in the **Common Name** field, and possibly several more in the **Subject Alternative Names** field. (CloudFront supports wildcard characters in certificate domain names.)

Important

When you add an alternate domain name to a distribution, CloudFront checks that the alternate domain name is covered by the certificate that you've attached. The certificate must cover the alternate domain name in the subject alternate name (SAN) field of the certificate. This means the SAN field must contain an exact match for the alternate domain name, or contain a wildcard at the same level of the alternate domain name that you're adding.

For more information, see [Requirements for Using Alternate Domain Names \(p. 78\)](#).

One of the domain names in the certificate must match the domain name that you specify for Origin Domain Name. If no domain name matches, CloudFront returns HTTP status code 502 (Bad Gateway) to the viewer.

Minimum SSL Protocol Version

If you're using dedicated IP addresses, set the minimum SSL protocol version for the connection between viewers and CloudFront by choosing a security policy.

For more information, see [Security Policy \(p. 56\)](#) in the topic [Values That You Specify When You Create or Update a Distribution \(p. 38\)](#).

Supported HTTP Versions

If you associate one certificate with more than one CloudFront distribution, all the distributions associated with the certificate must use the same option for [Supported HTTP Versions \(p. 57\)](#). You specify this option when you create or update a CloudFront distribution.

Quotas on Using SSL/TLS Certificates with CloudFront (HTTPS Between Viewers and CloudFront Only)

Note the following quotas (formerly known as limits) on using SSL/TLS certificates with CloudFront. These quotas apply only to the SSL/TLS certificates that you provision by using AWS Certificate Manager (ACM), that you import into ACM, or upload to the IAM certificate store for HTTPS communication between viewers and CloudFront.

Maximum Number of Certificates per CloudFront Distribution

You can associate a maximum of one SSL/TLS certificate with each CloudFront distribution.

Maximum Number of Certificates that You Can Import into ACM or Upload to the IAM Certificate Store

If you obtained your SSL/TLS certificates from a third-party CA, you must store the certificates in one of the following locations:

- **AWS Certificate Manager** – For the current quota on the number of ACM certificates, see [Quotas](#) in the *AWS Certificate Manager User Guide*. The listed quota is a total that includes certificates that you provision by using ACM and certificates that you import into ACM.
- **IAM certificate store** – For the current quota (formerly known as limit) on the number of certificates that you can upload to the IAM certificate store for an AWS account, see [IAM and STS Limits](#) in the *IAM User Guide*. You can [request a higher quota in the AWS Management Console](#).

Maximum Number of Certificates per AWS Account (Dedicated IP Addresses Only)

If you want to serve HTTPS requests by using dedicated IP addresses, note the following:

- By default, CloudFront gives you permission to use two certificates with your AWS account, one for everyday use and one for when you need to rotate certificates for multiple distributions.
- If you need more than two custom SSL/TLS certificates for your AWS account, go to the [Support Center](#) and create a case. Indicate how many certificates that you need permission to use, and describe the circumstances in your request. We'll update your account as soon as possible.

Using the Same Certificate for CloudFront Distributions that Were Created by Using Different AWS Accounts

If you're using a third-party CA and you want to use the same certificate with multiple CloudFront distributions that were created by using different AWS accounts, you must import the certificate into ACM or upload it to the IAM certificate store once for each AWS account.

If you're using certificates provided by ACM, you can't configure CloudFront to use certificates that were created by a different AWS account.

Using the Same Certificate for CloudFront and for Other AWS Services (IAM Certificate Store Only)

If you bought a certificate from a trusted certificate authority such as Comodo, DigiCert, or Symantec, you can use the same certificate for CloudFront and for other AWS services. If you're importing the certificate into ACM, you need to import it only once to use it for multiple AWS services.

If you're using certificates provided by ACM, the certificates are stored in ACM.

Using the Same Certificate for Multiple CloudFront Distributions

You can use the same certificate for any or all of the CloudFront distributions that you're using to serve HTTPS requests. Note the following:

- You can use the same certificate both for serving requests using dedicated IP addresses and for serving requests using SNI.
- You can associate only one certificate with each distribution.
- Each distribution must include one or more alternate domain names that also appear in the **Common Name** field or the **Subject Alternative Names** field in the certificate.
- If you're serving HTTPS requests using dedicated IP addresses and you created all of your distributions by using the same AWS account, you can significantly reduce your cost by using the same certificate for all distributions. CloudFront charges for each certificate, not for each distribution.

For example, suppose you create three distributions by using the same AWS account, and you use the same certificate for all three distributions. You would be charged only one fee for using dedicated IP addresses.

However, if you're serving HTTPS requests using dedicated IP addresses and using the same certificate to create CloudFront distributions in different AWS accounts, each account is charged

the fee for using dedicated IP addresses. For example, if you create three distributions by using three different AWS accounts and you use the same certificate for all three distributions, each account is charged the full fee for using dedicated IP addresses.

Configuring Alternate Domain Names and HTTPS

To use alternate domain names in the URLs for your files and to use HTTPS between viewers and CloudFront, perform the applicable procedures.

Topics

- [Requesting Permission to Use Three or More Dedicated IP SSL/TLS Certificates \(p. 139\)](#)
- [Getting an SSL/TLS Certificate \(p. 139\)](#)
- [Importing an SSL/TLS Certificate \(p. 140\)](#)
- [Updating Your CloudFront Distribution \(p. 140\)](#)

Requesting Permission to Use Three or More Dedicated IP SSL/TLS Certificates

If you need permission to permanently associate three or more SSL/TLS dedicated IP certificates with CloudFront, perform the following procedure. For more details about HTTPS requests, see [Choosing How CloudFront Serves HTTPS Requests \(p. 132\)](#).

Note

This procedure is for using three or more dedicated IP certificates across your CloudFront distributions. The default value is 2. Keep in mind you cannot bind more than one SSL certificate to a distribution.

You can only associate a single SSL/TLS certificate to a CloudFront distribution at a time. This number is for the total number of dedicated IP SSL certificates you can use across all of your CloudFront distributions.

To request permission to use three or more certificates with a CloudFront distribution

1. Go to the [Support Center](#) and create a case.
2. Indicate how many certificates you need permission to use, and describe the circumstances in your request. We'll update your account as soon as possible.
3. Continue with the next procedure.

Getting an SSL/TLS Certificate

Get an SSL/TLS certificate if you don't already have one. For more information, see the applicable documentation:

- To use a certificate provided by AWS Certificate Manager (ACM), see the [AWS Certificate Manager User Guide](#). Then skip to [Updating Your CloudFront Distribution \(p. 140\)](#).

Note

We recommend that you use ACM to provision, manage, and deploy SSL/TLS certificates on AWS managed resources. You must request an ACM certificate in the US East (N. Virginia) Region.

- To get a certificate from a third-party certificate authority (CA), see the documentation provided by the certificate authority. When you have the certificate, continue with the next procedure.
- To create a self-signed certificate, see the documentation for the application that you're using to create and sign the certificate. Then continue with the next procedure.

Importing an SSL/TLS Certificate

If you got your certificate from a third-party CA, import the certificate into ACM or upload it to the IAM certificate store:

ACM (Recommended)

ACM lets you import third-party certificates from the ACM console, as well as programmatically. For information about importing a certificate to ACM, see [Importing Certificates into AWS Certificate Manager](#) in the *AWS Certificate Manager User Guide*. You must import the certificate in the US East (N. Virginia) Region.

IAM certificate store

If ACM is not available in your Region, use the following AWS CLI command to upload your third-party certificate to the IAM certificate store. (For a list of the Regions where ACM is available, see [AWS Certificate Manager](#) in the "AWS Regions and Endpoints" chapter of the *Amazon Web Services General Reference*.)

```
aws iam upload-server-certificate --server-certificate-name CertificateName --  
certificate-body file://public_key_certificate_file --private-key file://privatekey.pem --  
--certificate-chain file://certificate_chain_file --path /cloudfront/path/
```

Note the following:

- **AWS account** – You must upload the certificate to the IAM certificate store using the same AWS account that you used to create your CloudFront distribution.
- **--path parameter** – When you upload the certificate to IAM, the value of the --path parameter (certificate path) must start with /cloudfront/, for example, /cloudfront/production/ or /cloudfront/test/. The path must end with a /.
- **Existing certificates** – You must specify values for the --server-certificate-name and --path parameters that are different from the values that are associated with existing certificates.
- **Using the CloudFront console** – The value that you specify for the --server-certificate-name parameter in the AWS CLI, for example, myServerCertificate, appears in the **SSL Certificate** list in the CloudFront console.
- **Using the CloudFront API** – Make note of the alphanumeric string that the AWS CLI returns, for example, AS1A2M3P4L5E67SIIXR3J. This is the value that you will specify in the IAMCertificateId element. You don't need the IAM ARN, which is also returned by the CLI.

For more information about the AWS CLI, see the [AWS Command Line Interface User Guide](#) and the [AWS CLI Command Reference](#).

Updating Your CloudFront Distribution

To update settings for your distribution, perform the following procedure:

To configure your CloudFront distribution for alternate domain names

1. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
2. Choose the ID for the distribution that you want to update.
3. On the **General** tab, choose **Edit**.
4. Update the following values:

Alternate Domain Names (CNAMEs)

Add the applicable alternate domain names. Separate domain names with commas, or type each domain name on a new line.

SSL Certificate (Web Distributions Only)

Choose **Custom SSL Certificate**, and choose a certificate from the list.

Up to 100 certificates are listed here. If you have more than 100 certificates and you don't see the certificate that you want to add, you can type a certificate ARN in the field to choose it.

If you uploaded a certificate to the IAM certificate store but it's not listed, and you can't choose it by typing the name in the field, review the procedure [Importing an SSL/TLS Certificate \(p. 140\)](#) to confirm that you correctly uploaded the certificate.

Important

After you associate your SSL/TLS certificate with your CloudFront distribution, do not delete the certificate from ACM or the IAM certificate store until you remove the certificate from all distributions and until the status of the distributions has changed to **Deployed**.

Clients Supported (Web Distributions Only)

Choose the applicable option:

- **All clients:** CloudFront serves your HTTPS content using dedicated IP addresses. If you select this option, you incur additional charges when you associate your SSL/TLS certificate with a distribution that is enabled. For more information, see [Amazon CloudFront Pricing](#).
- **Only clients that Support Server Name Indication (SNI):** Older browsers or other clients that don't support SNI must use another method to access your content.

For more information, see [Choosing How CloudFront Serves HTTPS Requests \(p. 132\)](#).

5. Choose **Yes, Edit**.
6. Configure CloudFront to require HTTPS between viewers and CloudFront:
 - a. On the **Behaviors** tab, choose the cache behavior that you want to update, and choose **Edit**.
 - b. Specify one of the following values for **Viewer Protocol Policy**:

Redirect HTTP to HTTPS

Viewers can use both protocols, but HTTP requests are automatically redirected to HTTPS requests. CloudFront returns HTTP status code 301 (**Moved Permanently**) along with the new HTTPS URL. The viewer then resubmits the request to CloudFront using the HTTPS URL.

Important

CloudFront doesn't redirect **DELETE**, **OPTIONS**, **PATCH**, **POST**, or **PUT** requests from HTTP to HTTPS. If you configure a cache behavior to redirect to HTTPS, CloudFront responds to HTTP **DELETE**, **OPTIONS**, **PATCH**, **POST**, or **PUT** requests for that cache behavior with HTTP status code 403 (**Forbidden**).

When a viewer makes an HTTP request that is redirected to an HTTPS request, CloudFront charges for both requests. For the HTTP request, the charge is only for the request and for the headers that CloudFront returns to the viewer. For the HTTPS request, the charge is for the request, and for the headers and the file returned by your origin.

HTTPS Only

Viewers can access your content only if they're using HTTPS. If a viewer sends an HTTP request instead of an HTTPS request, CloudFront returns HTTP status code 403 (Forbidden) and does not return the file.

- c. Choose **Yes, Edit**.
 - d. Repeat steps a through c for each additional cache behavior that you want to require HTTPS for between viewers and CloudFront.
7. Confirm the following before you use the updated configuration in a production environment:
- The path pattern in each cache behavior applies only to the requests that you want viewers to use HTTPS for.
 - The cache behaviors are listed in the order that you want CloudFront to evaluate them in. For more information, see [Path Pattern \(p. 47\)](#).
 - The cache behaviors are routing requests to the correct origins.

Determining the Size of the Public Key in an SSL/TLS Certificate

When you're using CloudFront alternate domain names and HTTPS, the maximum size of the public key in an SSL/TLS certificate is 2048 bits. (This is the key size, not the number of characters in the public key.) If you use AWS Certificate Manager for your certificates, although ACM supports larger keys, you cannot use the larger keys with CloudFront.

You can determine the size of the public key by running the following OpenSSL command:

```
openssl x509 -in path and filename of SSL/TLS certificate -text -noout
```

Where:

- `-in` specifies the path and file name of your SSL/TLS certificate.
- `-text` causes OpenSSL to display the length of the public key in bits.
- `-noout` prevents OpenSSL from displaying the public key.

Example output:

```
Public-Key: (2048 bit)
```

Increasing the Quotas for SSL/TLS Certificates

There are quotas (formerly known as limits) on the number of SSL/TLS certificates that you can import into [AWS Certificate Manager](#) or upload to [AWS Identity and Access Management](#). There also is a quota on the number of SSL/TLS certificates that you can use with an AWS account when you configure CloudFront to serve HTTPS requests by using dedicated IP addresses. However, you can request higher quotas.

Topics

- [Certificates That You Can Import into ACM \(p. 143\)](#)
- [Certificates That You Can Upload to IAM \(p. 143\)](#)
- [Certificates That You Can Use with Dedicated IP Addresses \(p. 143\)](#)

Certificates That You Can Import into ACM

For the quota on the number of certificates that you can import into ACM, see [Quotas](#) in the *AWS Certificate Manager User Guide*.

To request a higher quota, [create a case](#) in the AWS Support Center. Specify the following values:

- Accept the default value of **Service limit increase**.
- For **Limit type**, choose **Certificate Manager**.
- For **Region**, choose the AWS Region where you want to import certificates.
- For **Limit**, choose **Number of ACM certificates**.

Then fill out the rest of the form and submit it.

Certificates That You Can Upload to IAM

For the quota (formerly known as limit) on the number of certificates that you can upload to IAM, see [IAM and STS Limits](#) in the *IAM User Guide*.

To request a higher quota, [create a case](#) in the AWS Support Center. Specify the following values:

- Accept the default value of **Service limit increase**.
- For **Limit type**, choose **Certificate Manager**.
- For **Region**, choose the AWS Region where you want to import certificates.
- For **Limit**, choose **Server Certificate Limit (IAM)**.

Then fill out the rest of the form and submit it.

Certificates That You Can Use with Dedicated IP Addresses

For the quota (formerly known as limit) on the number of SSL certificates that you can use for each AWS account when serving HTTPS requests using dedicated IP addresses, see [Quotas on SSL Certificates \(Web Distributions Only\)](#) (p. 499).

To request a higher quota, [create a case](#) in the AWS Support Center. Specify the following values:

- Accept the default value of **Service limit increase**.
- For **Limit Type**, choose **CloudFront Distributions**.
- For **Limit**, choose **Dedicated IP SSL Certificate Limit per Account**.

Then fill out the rest of the form and submit it.

Rotating SSL/TLS Certificates

If you're using certificates provided by AWS Certificate Manager (ACM), you don't need to rotate SSL/TLS certificates. ACM manages certificate renewals for you. For more information, see [Managed Renewal](#) in the *AWS Certificate Manager User Guide*.

Note

ACM does not manage certificate renewals for certificates that you acquire from third-party certificate authorities and import into ACM.

If you're using a third-party certificate authority and you imported certificates into ACM (recommended) or uploaded them to the IAM certificate store, you must occasionally replace one certificate with another. For example, you must replace a certificate when the expiration date on the certificate approaches.

Important

If you configured CloudFront to serve HTTPS requests by using dedicated IP addresses, you might incur an additional, pro-rated charge for using one or more additional certificates while you're rotating certificates. We recommend that you update your distributions promptly to minimize the additional charge.

To rotate certificates, perform the following procedure. Viewers can continue to access your content while you rotate certificates as well as after the process is complete.

To rotate SSL/TLS certificates

1. [Increasing the Quotas for SSL/TLS Certificates \(p. 142\)](#) to determine whether you need permission to use more SSL certificates. If so, request permission and wait until permission is granted before you continue with step 2.
2. Import the new certificate into ACM or upload it to IAM. For more information, see [Importing an SSL/TLS Certificate](#) in the *Amazon CloudFront Developer Guide*.
3. Update your distributions one at a time to use the new certificate. For more information, see [Listing, Viewing, and Updating CloudFront Distributions](#) in the *Amazon CloudFront Developer Guide*.
4. (Optional) After you have updated all of your CloudFront distributions, you can delete the old certificate from ACM or from IAM.

Important

Do not delete an SSL/TLS certificate until you remove it from all distributions and until the status of the distributions that you have updated has changed to `Deployed`.

Reverting from a Custom SSL/TLS Certificate to the Default CloudFront Certificate

If you configured CloudFront to use HTTPS between viewers and CloudFront, and you configured CloudFront to use a custom SSL/TLS certificate, you can change your configuration to use the default CloudFront SSL/TLS certificate. The process depends on whether you've used your distribution to distribute your content:

- If you have not used your distribution to distribute your content, you can just change the configuration. For more information, see [Updating a Distribution \(p. 63\)](#).
- If you have used your distribution to distribute your content, you must create a new CloudFront distribution and change the URLs for your files to reduce or eliminate the amount of time that your content is unavailable. To do that, perform the following procedure.

To revert to the default CloudFront certificate

1. Create a new CloudFront distribution with the desired configuration. For **SSL Certificate**, choose **Default CloudFront Certificate (*.cloudfront.net)**.

For more information, see [Steps for Creating a Distribution \(Overview\) \(p. 36\)](#).

2. For files that you're distributing using CloudFront, update the URLs in your application to use the domain name that CloudFront assigned to the new distribution. For example, change `https://www.example.com/images/logo.png` to `https://d1111111abcdef8.cloudfront.net/images/logo.png`.

3. Either delete the distribution that is associated with a custom SSL/TLS certificate, or update the distribution to change the value of **SSL Certificate to Default CloudFront Certificate (*.cloudfront.net)**. For more information, see [Updating a Distribution \(p. 63\)](#).

Important

Until you complete this step, AWS continues to charge you for using a custom SSL/TLS certificate.

4. (Optional) Delete your custom SSL/TLS certificate.
 - a. Run the AWS CLI command `list-server-certificates` to get the certificate ID of the certificate that you want to delete. For more information, see [list-server-certificates](#) in the *AWS CLI Command Reference*.
 - b. Run the AWS CLI command `delete-signing-certificate` to delete the certificate. For more information, see [delete-signing-certificate](#) in the *AWS CLI Command Reference*.

Switching from a Custom SSL/TLS Certificate with Dedicated IP Addresses to SNI

If you configured CloudFront to use a custom SSL/TLS certificate with dedicated IP addresses, you can switch to using a custom SSL/TLS certificate with SNI instead and eliminate the charge that is associated with dedicated IP addresses. The following procedure shows you how.

Important

This update to your CloudFront configuration has no effect on viewers that support SNI. Viewers can access your content before and after the change, as well as while the change is propagating to CloudFront edge locations. Viewers that don't support SNI cannot access your content after the change. For more information, see [Choosing How CloudFront Serves HTTPS Requests \(p. 132\)](#).

To switch from a custom SSL/TLS certificate with dedicated IP addresses to SNI

1. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
2. Choose the ID of the distribution that you want to view or update.
3. Choose **Distribution Settings**.
4. On the **General** tab, choose **Edit**.
5. Change the setting of **Custom SSL Client Support** to **Only Clients that Support Server Name Indication (SNI)**.
6. Choose **Yes, Edit**.

Serving Private Content with Signed URLs and Signed Cookies

Many companies that distribute content over the internet want to restrict access to documents, business data, media streams, or content that is intended for selected users, for example, users who have paid a fee. To securely serve this private content by using CloudFront, you can do the following:

- Require that your users access your private content by using special CloudFront signed URLs or signed cookies.
- Require that your users access your content by using CloudFront URLs, not URLs that access content directly on the origin server (for example, Amazon S3 or a private HTTP server). Requiring CloudFront

URLs isn't necessary, but we recommend it to prevent users from bypassing the restrictions that you specify in signed URLs or signed cookies.


Topics

- [Overview of Serving Private Content \(p. 146\)](#)
- [Task List: Serving Private Content \(p. 149\)](#)
- [Specifying the AWS Accounts That Can Create Signed URLs and Signed Cookies \(Trusted Signers\) \(p. 150\)](#)
- [Choosing Between Signed URLs and Signed Cookies \(p. 156\)](#)
- [Using Signed URLs \(p. 157\)](#)
- [Using Signed Cookies \(p. 175\)](#)
- [Using a Linux Command and OpenSSL for Base64-Encoding and Encryption \(p. 189\)](#)
- [Code Examples for Creating a Signature for a Signed URL \(p. 189\)](#)

Overview of Serving Private Content

You can control user access to your private content in two ways, as shown in the following illustration:

1. Restrict access to files in CloudFront edge caches
2. Restrict access to files in your origin by doing one of the following:
 - Set up an origin access identity (OAI) for your Amazon S3 bucket
 - Configure custom headers for a private HTTP server (a custom origin)



Restricting Access to Files in CloudFront Edge Caches

You can configure CloudFront to require that users access your files using either *signed URLs* or *signed cookies*. You then develop your application either to create and distribute signed URLs to authenticated users or to send `Set-Cookie` headers that set signed cookies on the viewers for authenticated users. (To give a few users long-term access to a small number of files, you can also create signed URLs manually.)

When you create signed URLs or signed cookies to control access to your files, you can specify the following restrictions:

- An ending date and time, after which the URL is no longer valid.
- (Optional) The date and time that the URL becomes valid.
- (Optional) The IP address or range of addresses of the computers that can be used to access your content.

One part of a signed URL or a signed cookie is hashed and signed using the private key from a public/private key pair. When someone uses a signed URL or signed cookie to access a file, CloudFront compares the signed and unsigned portions of the URL or cookie. If they don't match, CloudFront doesn't serve the file.

You must use RSA-SHA1 for signing URLs or cookies. CloudFront doesn't accept other algorithms.

Restricting Access to Files in Amazon S3 Buckets

You can optionally secure the content in your Amazon S3 bucket so that users can access it through CloudFront but cannot access it directly by using Amazon S3 URLs. This prevents someone from

bypassing CloudFront and using the Amazon S3 URL to get content that you want to restrict access to. This step isn't required to use signed URLs, but we recommend it.

To require that users access your content through CloudFront URLs, you do the following tasks:

- Create a special CloudFront user called an *origin access identity* and associate it with your CloudFront distribution.
- Give the origin access identity permission to read the files in your bucket.
- Remove permission for anyone else to use Amazon S3 URLs to read the files.

For more information, see [Restricting Access to Amazon S3 Content by Using an Origin Access Identity \(p. 209\)](#).

Restricting Access to Files on Custom Origins

If you use a custom origin, you can optionally set up custom headers to restrict access. For CloudFront to get your files from a custom origin, the files must be publicly accessible. But by using custom headers, you can restrict access to your content so that users can access it only through CloudFront, not directly. This step isn't required to use signed URLs, but we recommend it.

To require that users access content through CloudFront, change the following settings in your CloudFront distributions:

Origin Custom Headers

Configure CloudFront to forward custom headers to your origin. See [Configuring CloudFront to Add Custom Headers to Origin Requests \(p. 283\)](#).

Viewer Protocol Policy

Configure your distribution to require viewers to use HTTPS to access CloudFront. See [Viewer Protocol Policy \(p. 49\)](#).

Origin Protocol Policy

Configure your distribution to require CloudFront to use the same protocol as viewers to forward requests to the origin. See [Origin Protocol Policy \(p. 45\)](#).

After you've made these changes, update your application on your custom origin to only accept requests that include these headers.

The combination of **Viewer Protocol Policy** and **Origin Protocol Policy** ensure that your custom headers are encrypted between the viewer and your origin. However, we recommend that you periodically do the following to rotate the custom headers that CloudFront forwards to your origin:

1. Update your CloudFront distribution to begin forwarding a new header to your custom origin.
2. Update your application to accept the new header as confirmation that the request is coming from CloudFront.
3. When viewer requests no longer include the header that you're replacing, update your application to no longer accept the old header as confirmation that the request is coming from CloudFront.

Task List: Serving Private Content

To configure CloudFront to serve private content, do the following tasks:

1. (Optional but recommended) Require your users to access your content only through CloudFront. The method that you use depends on whether you're using Amazon S3 or custom origins:

- **Amazon S3** – See [Restricting Access to Amazon S3 Content by Using an Origin Access Identity](#) (p. 209).
- **Custom origin** – See [Restricting Access to Files on Custom Origins](#) (p. 148).

Custom origins include Amazon EC2, Amazon S3 buckets configured as website endpoints, Elastic Load Balancing, and your own HTTP web servers.

2. Specify the AWS accounts that you want to use to create signed URLs or signed cookies. For more information, see [Specifying the AWS Accounts That Can Create Signed URLs and Signed Cookies \(Trusted Signers\)](#) (p. 150).
3. Write your application to respond to requests from authorized users either with signed URLs or with Set-Cookie headers that set signed cookies. Follow the steps in one of the following topics:
 - [Using Signed URLs](#) (p. 157)
 - [Using Signed Cookies](#) (p. 175)

If you're not sure which method to use, see [Choosing Between Signed URLs and Signed Cookies](#) (p. 156).

Specifying the AWS Accounts That Can Create Signed URLs and Signed Cookies (Trusted Signers)

Topics

- [Creating CloudFront Key Pairs for Your Trusted Signers \(p. 151\)](#)
- [Reformatting the CloudFront Private Key \(.NET and Java Only\) \(p. 152\)](#)
- [Adding Trusted Signers to Your Distribution \(p. 152\)](#)
- [Verifying that Trusted Signers Are Active \(Optional\) \(p. 154\)](#)
- [Rotating CloudFront Key Pairs \(p. 155\)](#)

To create signed URLs or signed cookies, you need at least one AWS account that has an active CloudFront key pair. This account is known as a trusted signer. The trusted signer has two purposes:

- As soon as you add the AWS account ID for your trusted signer to your distribution, CloudFront starts to require that users use signed URLs or signed cookies to access your files.
- When you create signed URLs or signed cookies, you use the private key from the trusted signer's key pair to sign a portion of the URL or the cookie. When someone requests a restricted file, CloudFront compares the signed portion of the URL or cookie with the unsigned portion to verify that the URL or cookie hasn't been tampered with. CloudFront also verifies that the URL or cookie is valid, meaning, for example, that the expiration date and time hasn't passed.

When you specify trusted signers, you also indirectly specify the files that require signed URLs or signed cookies:

- **Web distributions** – You add trusted signers to cache behaviors. If your distribution has only one cache behavior, users must use signed URLs or signed cookies to access any file associated with the distribution. If you create multiple cache behaviors and add trusted signers to some cache behaviors and not to others, you can require that users use signed URLs or signed cookies to access some files and not others.
- **RTMP distributions (signed URLs only)** – You add trusted signers to a distribution. After you add trusted signers to an RTMP distribution, users must use signed URLs to access any file associated with the distribution.

Note

To specify trusted signers for a distribution, you must use the CloudFront console or CloudFront API version 2009-09-09 or later.

To specify the accounts that are allowed to create signed URLs or signed cookies and to add the accounts to your CloudFront distribution, do the following tasks:

1. Decide which AWS accounts you want to use as trusted signers. Most CloudFront customers use the account that they used to create the distribution.
2. For each of the accounts that you selected in Step 1, create a CloudFront key pair. For more information, see [Creating CloudFront Key Pairs for Your Trusted Signers \(p. 151\)](#).
3. If you're using .NET or Java to create signed URLs or signed cookies, reformat the CloudFront private key. For more information, see [Reformatting the CloudFront Private Key \(.NET and Java Only\) \(p. 152\)](#).
4. In the distribution for which you're creating signed URLs or signed cookies, specify the AWS account IDs of your trusted signers. For more information, see [Adding Trusted Signers to Your Distribution \(p. 152\)](#).

5. (Optional) Verify that CloudFront recognizes that your trusted signers have active CloudFront key pairs. For more information, see [Verifying that Trusted Signers Are Active \(Optional\)](#) (p. 154).

Creating CloudFront Key Pairs for Your Trusted Signers

Each of the AWS accounts that you use to create CloudFront signed URLs or signed cookies—your trusted signers—must have its own CloudFront key pair, and the key pair must be active. Note that you can't substitute an Amazon EC2 key pair for a CloudFront key pair. When you create a CloudFront signed URL or signed cookie, you include the key pair ID for the trusted signer's key pair in the URL. Amazon EC2 does not make key pair IDs available.

To help secure your applications, we recommend that you change CloudFront key pairs every 90 days or more often. For more information, see [Rotating CloudFront Key Pairs](#) (p. 155).

You can create a key pair in the following ways:

- Create a key pair in the AWS Management Console and download the private key. See the procedure [To create CloudFront key pairs in the AWS Management Console](#) (p. 151).
- Create an RSA key pair by using an application such as OpenSSL, and upload the public key to the AWS Management Console. See the procedure [To create an RSA key pair and upload the public key in the AWS Management Console](#) (p. 151).

To create CloudFront key pairs in the AWS Management Console

1. Sign in to the AWS Management Console using the root credentials for an AWS account.

Important

IAM users can't create CloudFront key pairs. You must log in using root credentials to create key pairs.

2. On the *account-name* menu, click **Security Credentials**.
3. Expand **CloudFront Key Pairs**.
4. Confirm that you have no more than one active key pair. You can't create a key pair if you already have two active key pairs.
5. Click **Create New Key Pair**.
6. In the **Create Key Pair** dialog box, click **Download Private Key File**.
7. In the **Opening <filename>** dialog box, accept the default value of **Save File**, and click **OK** to download and save the private key for your CloudFront key pair.

Important

Save the private key for your CloudFront key pair in a secure location, and set permissions on the file so that only the desired administrator users can read it. If someone gets your private key, they can generate valid signed URLs and signed cookies and download your content. You cannot get the private key again, so if you lose or delete it, you must create a new CloudFront key pair.

8. Record the key pair ID for your key pair. (In the AWS Management Console, this is called the access key ID.) You'll use it when you create signed URLs or signed cookies.

To create an RSA key pair and upload the public key in the AWS Management Console

1. Use OpenSSL or another tool to create a key pair.

For example, if you're using OpenSSL, you can use the following command to generate a key pair with a length of 4096 bits and save it in the file `private_key.pem`:


```
openssl genrsa -out private_key.pem 4096
```

The resulting file contains both the public and the private key. To extract the public key from that file, run the following command:

```
openssl rsa -pubout -in private_key.pem -out public_key.pem
```

The public key (`public_key.pem`) is the file that you upload later in this procedure.

Note the following requirements for the key:

- The key pair must be an SSH-2 RSA key pair.
 - The key pair must be in base64 encoded PEM format.
 - The supported key lengths are 1024, 2048, and 4096 bits.
2. Sign in to the AWS Management Console using the root credentials for an AWS account.
Important
IAM users can't create CloudFront key pairs. You must log in using root credentials to create key pairs.
 3. On the *account-name* menu, click **Security Credentials**.
 4. Expand **CloudFront Key Pairs**.
 5. Confirm that you have no more than one active key pair. You can't upload your own key pair if you already have two active key pairs.
 6. Click **Upload Your Own Key Pair**.
 7. In the **Upload Your Own Key Pair** dialog box, click **Choose File** and choose the public key file that you created in step 1.
 8. Click **Upload**.

The **Upload Key Pair** dialog box clears, and the new key pair appears at the top of the list of CloudFront key pairs.
 9. Record the key pair ID for your key pair. (In the AWS Management Console, this is called the access key ID.) You'll use it when you create signed URLs or signed cookies.

Reformatting the CloudFront Private Key (.NET and Java Only)

If you're using .NET or Java to create signed URLs or signed cookies, you cannot use the private key from your key pair in the default .pem format to create the signature:

- **.NET framework** – Convert the private key to the XML format that the .NET framework uses. Several tools are available.
- **Java** – Convert the private key to DER format. To do this, you can use OpenSSL:

```
$ openssl pkcs8 -topk8 -nocrypt -in origin.pem -inform PEM -out new.der -  
outform DER
```

To ensure that the encoder works correctly, add the jar for the Bouncy Castle Java cryptography APIs to your project and then add the Bouncy Castle provider.

Adding Trusted Signers to Your Distribution

Trusted signers are the AWS accounts that can create signed URLs and signed cookies for a distribution. By default, no account, not even the account that created the distribution, is allowed to create signed

URLs or signed cookies. To specify the AWS accounts that you want to use as trusted signers, add the accounts to your distribution:

- **Web distributions** – Trusted signers are associated with cache behaviors. This allows you to require signed URLs or signed cookies for some files and not for others in the same distribution. Trusted signers can only create signed URLs or cookies for files that are associated with the corresponding cache behaviors. For example, if you have one trusted signer for one cache behavior and a different trusted signer for a different cache behavior, neither trusted signer can create signed URLs or cookies for files that are associated with the other cache behavior.
- **RTMP distributions (signed URLs only)** – Trusted signers are associated with the distribution. After you add trusted signers to an RTMP distribution, users must use signed URLs or signed cookies to access any of the objects associated with the distribution.

Important

Define path patterns and their sequence carefully so you don't either give users unintended access to your content or prevent them from accessing content that you want to be available to everyone. For example, suppose a request matches the path pattern for two cache behaviors. The first cache behavior does not require signed URLs or signed cookies and the second cache behavior does. Users will be able to access the files without using signed URLs or signed cookies because CloudFront processes the cache behavior that is associated with the first match.

For more information about path patterns, see [Path Pattern \(p. 47\)](#).

Important

If you're updating a distribution that you're already using to distribute content, add trusted signers only when you're ready to start generating signed URLs or signed cookies for your files, or CloudFront will reject the requests:

- **Web distributions** – After you add trusted signers to a cache behavior for a web distribution, users must use signed URLs or signed cookies to access the files that are associated with the cache behavior.
- **RTMP distributions (signed URLs only)** – After you add trusted signers to an RTMP distribution, users must use signed URLs to access any of the files associated with the distribution.

The maximum number of trusted signers depends on the type of distribution:

- **Web distributions** – A maximum of five for each cache behavior
- **RTMP distributions** – A maximum of five for the distribution

You can add trusted signers to your distribution using either the CloudFront console or the CloudFront API. See the following topic:

- [Adding Trusted Signers to Your Distribution Using the CloudFront Console \(p. 153\)](#)
- [Adding Trusted Signers to Your Distribution Using the CloudFront API \(p. 154\)](#)

Adding Trusted Signers to Your Distribution Using the CloudFront Console

To add trusted signers to your distribution using the CloudFront console

1. If you want to use only the AWS account that created the distribution as a trusted signer, skip to Step 2.

If you want to use other AWS accounts, get the AWS account ID for each account:

- a. Sign in to the AWS Management Console at <https://console.aws.amazon.com/console/home> using an account that you want to use as a trusted signer.
- b. In the upper-right corner of the console, click the name associated with the account, and click **My Account**.
- c. Under **Account Settings**, make note of the account ID.
- d. Sign out of the AWS Management Console.
- e. Repeat steps a through d for the other accounts that you want to use as trusted signers.
2. Open the Amazon CloudFront console at <https://console.aws.amazon.com/cloudfront/>, and sign in using the account that you used to create the distribution that you want to add trusted signers to.
3. Click the distribution ID.
4. Change to edit mode:
 - **Web distributions** – Click the **Behaviors** tab, click the behavior that you want to edit, and click **Edit**.
 - **RTMP distributions** – Click **Edit**.
5. For **Restrict Viewer Access (Use Signed URLs or Signed Cookies)**, click **Yes**.
6. For **Trusted Signers**, select the check boxes for your scenario:
 - **Self** – Select this check box if you want to use the current account (the account that you used to create the distribution).
 - **Specify Accounts** – Select this check box if you want to use other AWS accounts.
7. If you selected the **Specify Accounts** check box, enter AWS account IDs in the **AWS Account Number** field. These are the account IDs that you got in the first step of this procedure. Enter one account ID per line.
8. Click **Yes, Edit**.
9. If you're adding trusted signers to a web distribution and you have more than one cache behavior, repeat steps 4 through 8.

Adding Trusted Signers to Your Distribution Using the CloudFront API

You can use the CloudFront API to add the AWS account IDs for trusted signers to an existing distribution or to create a new distribution that includes trusted signers. In either case, specify the values in the `TrustedSigners` element. For web distributions, add the `TrustedSigners` element to one or more cache behaviors. For RTMP distributions, add the `TrustedSigners` element to the distribution.

See the following topics in the *Amazon CloudFront API Reference*:

- **Create a new web distribution** – [CreateDistribution](#)
- **Update an existing web distribution** – [UpdateDistribution](#)
- **Create a new RTMP distribution** – [CreateStreamingDistribution](#)
- **Update an existing RTMP distribution** – [UpdateStreamingDistribution](#)

Verifying that Trusted Signers Are Active (Optional)

After you add trusted signers to your distribution, you might want to verify that the signers are active. For a trusted signer to be active, the following must be true:

- The AWS account must have at least one active key pair. If you're rotating key pairs, the account will temporarily have two active key pairs, the old key pair and the new one.
- CloudFront must be aware of the active key pair. After you create a key pair, there can be a short period of time before CloudFront is aware that the key pair exists.

Note

To display a list of active trusted signers for a distribution, you currently must use the CloudFront API. A list of active trusted signers is not available in the CloudFront console.

Verifying that Trusted Signers Are Active Using the CloudFront API

To determine which trusted signers have active key pairs (are active trusted signers), you get the distribution and review the values in the `ActiveTrustedSigners` element. This element lists the AWS account ID of each account that the distribution identifies as a trusted signer. If the trusted signer has one or more active CloudFront key pairs, the `ActiveTrustedSigners` element also lists the key pair IDs. For more information, see the following topics in the *Amazon CloudFront API Reference*:

- **Web distributions** – [GetDistribution](#)
- **RTMP distributions** – [GetStreamingDistribution](#)

Rotating CloudFront Key Pairs

AWS recommends that you rotate (change) your active CloudFront key pairs every 90 days. To rotate CloudFront key pairs that you're using to create signed URLs or signed cookies without invalidating URLs or cookies that haven't expired yet, do the following tasks:

1. Create a new key pair for each of the accounts that you're using to create signed URLs. For more information, see [Creating CloudFront Key Pairs for Your Trusted Signers \(p. 151\)](#).
2. Verify that CloudFront is aware of the new key pairs. For more information, see [Verifying that Trusted Signers Are Active \(Optional\) \(p. 154\)](#).
3. Update your application to create signatures using the private keys from the new key pairs.
4. Confirm that URLs or cookies that you're signing using the new private keys are working.
5. Wait until the expiration date has passed in URLs or cookies that were signed using the old CloudFront key pairs.
6. Change the old CloudFront key pairs to **Inactive**:
 - a. Sign in to the AWS Management Console using the root credentials for an AWS account for which you want to make key pairs inactive.
 - b. On the *account-name* menu, click **Security Credentials**.
 - c. Expand **CloudFront Key Pairs**.
 - d. Choose specific key pairs, and then choose **Make Inactive**.
 - e. Repeat steps a through d for each of the AWS accounts for which you want to make key pairs inactive.
7. Reconfirm that URLs or cookies that you're signing using the new private keys are working.
8. Delete the old CloudFront key pairs:
 - a. Go to the [Your Security Credentials](#) page.
 - b. Expand **CloudFront Key Pairs**.
 - c. Choose specific key pairs, and then choose **Delete**.
9. Delete the old private keys from the location where you stored them.

Choosing Between Signed URLs and Signed Cookies

CloudFront signed URLs and signed cookies provide the same basic functionality: they allow you to control who can access your content. If you want to serve private content through CloudFront and you're trying to decide whether to use signed URLs or signed cookies, consider the following.

Use signed URLs in the following cases:

- You want to use an RTMP distribution. Signed cookies aren't supported for RTMP distributions.
- You want to restrict access to individual files, for example, an installation download for your application.
- Your users are using a client (for example, a custom HTTP client) that doesn't support cookies.

Use signed cookies in the following cases:

- You want to provide access to multiple restricted files, for example, all of the files for a video in HLS format or all of the files in the subscribers' area of website.
- You don't want to change your current URLs.

If you are not currently using signed URLs, and if your (unsigned) URLs contain any of the following query string parameters, you cannot use either signed URLs or signed cookies:

- Expires
- Policy
- Signature
- Key-Pair-Id

CloudFront assumes that URLs that contain any of those query string parameters are signed URLs, and therefore won't look at signed cookies.

Using Both Signed URLs and Signed Cookies

If you use both signed URLs and signed cookies to control access to the same files and a viewer uses a signed URL to request a file, CloudFront determines whether to return the file to the viewer based only on the signed URL.

Using Signed URLs

Topics

- [Choosing Between Canned and Custom Policies for Signed URLs \(p. 157\)](#)
- [How Signed URLs Work \(p. 158\)](#)
- [Choosing How Long Signed URLs Are Valid \(p. 158\)](#)
- [When Does CloudFront Check the Expiration Date and Time in a Signed URL? \(p. 159\)](#)
- [Sample Code and Third-Party Tools \(p. 159\)](#)
- [Creating a Signed URL Using a Canned Policy \(p. 160\)](#)
- [Creating a Signed URL Using a Custom Policy \(p. 166\)](#)

A signed URL includes additional information, for example, an expiration date and time, that gives you more control over access to your content. This additional information appears in a policy statement, which is based on either a canned policy or a custom policy. The differences between canned and custom policies are explained in the next two sections.

Note

You can create some signed URLs using canned policies and create some signed URLs using custom policies for the same distribution.

Choosing Between Canned and Custom Policies for Signed URLs

When you create a signed URL, you write a policy statement in JSON format that specifies the restrictions on the signed URL, for example, how long the URL is valid. You can use either a canned policy or a custom policy. Here's how canned and custom policies compare:

Description	Canned Policy	Custom Policy
You can reuse the policy statement for multiple files. To reuse the policy statement, you must use wildcard characters in the <code>Resource</code> object. For more information, see Values that You Specify in the Policy Statement for a Signed URL That Uses a Custom Policy (p. 169) .	No	Yes
You can specify the date and time that users can begin to access your content.	No	Yes (optional)
You can specify the date and time that users can no longer access your content.	Yes	Yes
You can specify the IP address or range of IP addresses of the users who can access your content.	No	Yes (optional)
The signed URL includes a base64-encoded version of the policy, which results in a longer URL.	No	Yes

For information about creating signed URLs using a *canned* policy, see [Creating a Signed URL Using a Canned Policy \(p. 160\)](#).

For information about creating signed URLs using a *custom* policy, see [Creating a Signed URL Using a Custom Policy \(p. 166\)](#).

How Signed URLs Work

Here's an overview of how you configure CloudFront and Amazon S3 for signed URLs and how CloudFront responds when a user uses a signed URL to request a file.

1. In your CloudFront distribution, specify one or more trusted signers, which are the AWS accounts that you want to have permission to create signed URLs.

For more information, see [Specifying the AWS Accounts That Can Create Signed URLs and Signed Cookies \(Trusted Signers\)](#) (p. 150).

2. You develop your application to determine whether a user should have access to your content and to create signed URLs for the files or parts of your application that you want to restrict access to. For more information, see the following topics:

- [Creating a Signed URL Using a Canned Policy](#) (p. 160)
- [Creating a Signed URL Using a Custom Policy](#) (p. 166)

3. A user requests a file for which you want to require signed URLs.
4. Your application verifies that the user is entitled to access the file: they've signed in, they've paid for access to the content, or they've met some other requirement for access.
5. Your application creates and returns a signed URL to the user.
6. The signed URL allows the user to download or stream the content.

This step is automatic; the user usually doesn't have to do anything additional to access the content. For example, if a user is accessing your content in a web browser, your application returns the signed URL to the browser. The browser immediately uses the signed URL to access the file in the CloudFront edge cache without any intervention from the user.

7. CloudFront uses the public key to validate the signature and confirm that the URL hasn't been tampered with. If the signature is invalid, the request is rejected.

If the signature is valid, CloudFront looks at the policy statement in the URL (or constructs one if you're using a canned policy) to confirm that the request is still valid. For example, if you specified a beginning and ending date and time for the URL, CloudFront confirms that the user is trying to access your content during the time period that you want to allow access.

If the request meets the requirements in the policy statement, CloudFront does the standard operations: determines whether the file is already in the edge cache, forwards the request to the origin if necessary, and returns the file to the user.

Note

Signed CloudFront URLs cannot contain extra query string arguments. If you add a query string to a signed URL after you create it, the URL returns an HTTP 403 status.

Choosing How Long Signed URLs Are Valid

You can distribute private content using a signed URL that is valid for only a short time—possibly for as little as a few minutes. Signed URLs that are valid for such a short period are good for distributing content on-the-fly to a user for a specific purpose, such as distributing movie rentals or music downloads to customers on demand. If your signed URLs will be valid for just a short period, you'll probably want to generate them automatically using an application that you develop. When the user starts to download a file or starts to play a media file, CloudFront compares the expiration time in the URL with the current time to determine whether the URL is still valid.

You can also distribute private content using a signed URL that is valid for a longer time, possibly for years. Signed URLs that are valid for a longer period are useful for distributing private content to known

users, such as distributing a business plan to investors or distributing training materials to employees. You can develop an application to generate these longer-term signed URLs for you.

When Does CloudFront Check the Expiration Date and Time in a Signed URL?

When CloudFront checks the expiration date and time in a signed URL to determine whether the URL is still valid depends on whether the URL is for a web distribution or an RTMP distribution:

- **Web distributions** – CloudFront checks the expiration date and time in a signed URL at the time of the HTTP request. If a client begins to download a large file immediately before the expiration time, the download should complete even if the expiration time passes during the download. If the TCP connection drops and the client tries to restart the download after the expiration time passes, the download will fail.

If a client uses Range GETs to get a file in smaller pieces, any GET request that occurs after the expiration time passes will fail. For more information about Range GETs, see [How CloudFront Processes Partial Requests for an Object \(Range GETs\)](#) (p. 285).

- **RTMP distributions** – CloudFront checks the expiration time in a signed URL at the start of a play event. If a client starts to play a media file before the expiration time passes, CloudFront allows the entire media file to play. However, depending on the media player, pausing and restarting might trigger another play event. Skipping to another position in the media file will trigger another play event. If the subsequent play event occurs after the expiration time passes, CloudFront won't serve the media file.

Note

[Adobe designated Flash as end-of-life at the end of 2020](#). As a result, Amazon CloudFront will no longer support Adobe Flash Media Server and will be deprecating Real-Time Messaging Protocol (RTMP) distributions by December 31, 2020. For more information, [read the full announcement on the Amazon CloudFront discussion forum](#).

Sample Code and Third-Party Tools

For sample code that creates the hashed and signed part of signed URLs, see the following topics:

- [Create a URL Signature Using Perl](#) (p. 190)
- [Create a URL Signature Using PHP](#) (p. 198)
- [Create a URL Signature Using C# and the .NET Framework](#) (p. 203)
- [Create a URL Signature Using Java](#) (p. 208)

Creating a Signed URL Using a Canned Policy

To create a signed URL using a canned policy, do the following procedure.

To create a signed URL using a canned policy

1. If you're using .NET or Java to create signed URLs, and if you haven't reformatted the private key for your key pair from the default .pem format to a format compatible with .NET or with Java, do so now. For more information, see [Reformatting the CloudFront Private Key \(.NET and Java Only\)](#) (p. 152).
2. Concatenate the following values in the specified order, and remove the whitespace (including tabs and newline characters) between the parts. You might have to include escape characters in the string in application code. All values have a type of String. Each part is keyed by number (1) to the two examples that follow.

1 Base URL for the file

The base URL is the CloudFront URL that you would use to access the file if you were not using signed URLs, including your own query string parameters, if any. For more information about the format of URLs for web distributions, see [Customizing the URL Format for Files in CloudFront](#) (p. 106).

The following examples show values that you specify for web distributions.

- The following CloudFront URL is for a file in a web distribution (using the CloudFront domain name). Note that `image.jpg` is in an `images` directory. The path to the file in the URL must match the path to the file on your HTTP server or in your Amazon S3 bucket.

```
http://d1111111abcdef8.cloudfront.net/images/image.jpg
```

- The following CloudFront URL includes a query string:

```
http://d1111111abcdef8.cloudfront.net/images/image.jpg?size=large
```

- The following CloudFront URLs are for files in a web distribution. Both use an alternate domain name; the second one includes a query string:

```
http://www.example.com/images/image.jpg
```

```
http://www.example.com/images/image.jpg?color=red
```

- The following CloudFront URL is for files in a web distribution that uses an alternate domain name and the HTTPS protocol:

```
https://www.example.com/images/image.jpg
```

For RTMP distributions, the following examples are for files in two different video formats, MP4 and FLV:

- **MP4** – `mp4:sydney-vacation.mp4`
- **FLV** – `sydney-vacation`
- **FLV** – `sydney-vacation.flv`

Note

For .flv files, whether you include the `.flv` filename extension depends on your player. To serve MP3 audio files or H.264/MPEG-4 video files, you might need to prefix the file name with `mp3:` or `mp4:`. Some media players can be configured to add the prefix automatically. The media player might also require you to specify the file name without the file extension (for example, `sydney-vacation` instead of `sydney-vacation.mp4`).

2 ?

The ? indicates that query string parameters follow the base URL. Include the ? even if you don't have any query string parameters of your own.

3 *Your query string parameters, if any&*

This value is optional. If you want to add your own query string parameters, for example:

`color=red&size=medium`

then add the parameters after the ? (see **2**) and before the Expires parameter. In certain rare circumstances, you might need to put your query string parameters after Key-Pair-Id.

Important

Your parameters cannot be named Expires, Signature, or Key-Pair-Id.

If you add your own parameters, append an & after each one, including the last one.

4 *Expires=date and time in Unix time format (in seconds) and Coordinated Universal Time (UTC)*

The date and time that you want the URL to stop allowing access to the file.

Specify the expiration date and time in Unix time format (in seconds) and Coordinated Universal Time (UTC). For example, January 1, 2013 10:00 am UTC converts to 1357034400 in Unix time format. To use epoch time, use a 32-bit integer for a date which can be no later than 2147483647 (January 19th, 2038 at 03:14:07 UTC). For information about UTC, see *RFC 3339, Date and Time on the Internet: Timestamps*, <http://tools.ietf.org/html/rfc3339>.

5 *&Signature=hashed and signed version of the policy statement*

A hashed, signed, and base64-encoded version of the JSON policy statement. For more information, see [Creating a Signature for a Signed URL That Uses a Canned Policy](#) (p. 162).

6 *&Key-Pair-Id=active CloudFront key pair Id for the key pair that you're using to generate the signature*

The ID for an active CloudFront key pair, for example, APKA9ONS7QCOWEXAMPLE. The CloudFront key pair ID tells CloudFront which public key to use to validate the signed URL. CloudFront compares the information in the signature with the information in the policy statement to verify that the URL has not been tampered with.

The key pair ID that you include in CloudFront signed URLs must be the ID of an active key pair for one of your trusted signers:

- **Web distributions** – The key pair must be associated with an AWS account that is one of the trusted signers for the cache behavior.
- **RTMP distributions** – The key pair must be associated with an AWS account that is one of the trusted signers for the distribution.

For more information, see [Specifying the AWS Accounts That Can Create Signed URLs and Signed Cookies \(Trusted Signers\)](#) (p. 150).

If you make a key pair inactive while rotating CloudFront key pairs, and if you're generating signed URLs programmatically, you must update your application to use a new active key pair for one of your trusted signers. If you're generating signed URLs manually, you must create new signed URLs. For more information about rotating key pairs, see [Rotating CloudFront Key Pairs](#) (p. 155).

Example signed URL for a web distribution:

1 `http://d111111abcdef8.cloudfront.net/image.jpg` **2**
3 `? color=red&size=medium&` **4** `Expires=1357034400`
5 `&Signature=nitfHRCrtziwO2HwPfWw-yYDhUF5EwRunQA-`
`j19DzZrvDh6hQ73lDx--ar3UocvvRQVw6EkC-GdpGQyyOSKQim-`
`TxAnW7d8F5Kkai9HVxOFIu-5jcQb0UEmatEXAMPLE3ReXySpLSMj0yCd3ZAB4UcBCAqEijkytL6f3fVYNGQI6`
6 `&Key-Pair-Id=APKA9ONS7QCOWEXAMPLE`

Example signed URL for an RTMP distribution:

1 `videos/mediafile.flv` **2** `?` **3** `color=red&size=medium&` **4**
`Expires=1357034400` **5** `&Signature=nitfHRCrtziwO2HwPfWw-yYDhUF5EwRunQA-`
`j19DzZrvDh6hQ73lDx--ar3UocvvRQVw6EkC-GdpGQyyOSKQim-`
`TxAnW7d8F5Kkai9HVxOFIu-5jcQb0UEmatEXAMPLE3ReXySpLSMj0yCd3ZAB4UcBCAqEijkytL6f3fVYNGQI6`
6 `&Key-Pair-Id=APKA9ONS7QCOWEXAMPLE`

Creating a Signature for a Signed URL That Uses a Canned Policy

To create the signature for a signed URL that uses a canned policy, you do the following procedures:

1. Create a policy statement. See [Creating a Policy Statement for a Signed URL That Uses a Canned Policy](#) (p. 162).
2. Sign the policy statement to create a signature. See [Creating a Signature for a Signed URL That Uses a Canned Policy](#) (p. 164).

Creating a Policy Statement for a Signed URL That Uses a Canned Policy

When you create a signed URL using a canned policy, the `Signature` parameter is a hashed and signed version of a policy statement. For signed URLs that use a canned policy, you don't include the policy statement in the URL, as you do for signed URLs that use a custom policy. To create the policy statement, do the following procedure.

To create the policy statement for a signed URL that uses a canned policy

1. Construct the policy statement using the following JSON format and using UTF-8 character encoding. Include all punctuation and other literal values exactly as specified. For information about the `Resource` and `DateLessThan` parameters, see [Values that You Specify in the Policy Statement for a Signed URL That Uses a Canned Policy](#) (p. 163).

```
{
  "Statement": [
    {
      "Resource": "base URL or stream name",
      "Condition": {
        "DateLessThan": {
          "AWS:EpochTime": ending date and time in Unix time format and UTC
        }
      }
    }
  ]
}
```

2. Remove all whitespace (including tabs and newline characters) from the policy statement. You might have to include escape characters in the string in application code.

Values that You Specify in the Policy Statement for a Signed URL That Uses a Canned Policy

When you create a policy statement for a canned policy, you specify the following values.

Resource

The value that you specify depends on whether you're creating the signed URL for a web distribution or an RTMP distribution.

Note

You can specify only one value for Resource.

Web distributions

The base URL including your query strings, if any, but excluding the CloudFront `Expires`, `Signature`, and `Key-Pair-Id` parameters, for example:

```
http://d111111abcdef8.cloudfront.net/images/horizon.jpg?
size=large&license=yes
```

Note the following:

- **Protocol** – The value must begin with `http://` or `https://`.
- **Query string parameters** – If you have no query string parameters, omit the question mark.
- **Alternate domain names** – If you specify an alternate domain name (CNAME) in the URL, you must specify the alternate domain name when referencing the file in your web page or application. Do not specify the Amazon S3 URL for the object.

RTMP distributions

Include only the stream name. For example, if the full URL for a streaming video is:

```
rtmp://s5c39gqb8ow64r.cloudfront.net/videos/cfx/st/mp3_name.mp3
```

then use the following value for Resource:

```
videos/mp3_name
```

Do not include a prefix such as `mp3:` or `mp4:`. Also, depending on the player you're using, you might have to omit the file extension from the value of Resource. For example, you might need to use `sydney-vacation` instead of `sydney-vacation.flv`.

DateLessThan

The expiration date and time for the URL in Unix time format (in seconds) and Coordinated Universal Time (UTC). For example, January 1, 2013 10:00 am UTC converts to 1357034400 in Unix time format.

This value must match the value of the `Expires` query string parameter in the signed URL. Do not enclose the value in quotation marks.

For more information, see [When Does CloudFront Check the Expiration Date and Time in a Signed URL?](#) (p. 159).

Example Policy Statement for a Signed URL That Uses a Canned Policy

When you use the following example policy statement in a signed URL, a user can access the file `http://d111111abcdef8.cloudfront.net/horizon.jpg` until January 1, 2013 10:00 am UTC:

```
{
  "Statement": [
    {
      "Resource": "http://d111111abcdef8.cloudfront.net/horizon.jpg?
size=large&license=yes",
```

```
        "Condition":{
            "DateLessThan":{
                "AWS:EpochTime":1357034400
            }
        }
    }
]
```

Creating a Signature for a Signed URL That Uses a Canned Policy

To create the value for the `Signature` parameter in a signed URL, you hash and sign the policy statement that you created in [Creating a Policy Statement for a Signed URL That Uses a Canned Policy](#) (p. 162). There are two versions of this procedure. Follow the procedure for your scenario:

- [Option 1: To create a signature for a web distribution or for an RTMP distribution \(without Adobe Flash Player\) by using a canned policy](#) (p. 164)
- [Option 2: To create a signature for an RTMP distribution by using a canned policy \(Adobe Flash Player\)](#) (p. 165)

For additional information and examples of how to hash, sign, and encode the policy statement, see:

- [Using a Linux Command and OpenSSL for Base64-Encoding and Encryption](#) (p. 189)
- [Code Examples for Creating a Signature for a Signed URL](#) (p. 189)

Option 1: To create a signature for a web distribution or for an RTMP distribution (without Adobe Flash Player) by using a canned policy

1. Use the SHA-1 hash function and RSA to hash and sign the policy statement that you created in the procedure [To create the policy statement for a signed URL that uses a canned policy](#) (p. 162). Use the version of the policy statement that no longer includes whitespace.

For the private key that is required by the hash function, use the private key that is associated with the active trusted signer.

Note

The method that you use to hash and sign the policy statement depends on your programming language and platform. For sample code, see [Code Examples for Creating a Signature for a Signed URL](#) (p. 189).

2. Remove whitespace (including tabs and newline characters) from the hashed and signed string.
3. Base64-encode the string using MIME base64 encoding. For more information, see [Section 6.8, Base64 Content-Transfer-Encoding](#) in *RFC 2045, MIME (Multipurpose Internet Mail Extensions) Part One: Format of Internet Message Bodies*.
4. Replace characters that are invalid in a URL query string with characters that are valid. The following table lists invalid and valid characters.

Replace these invalid characters	With these valid characters
+	- (hyphen)
=	_ (underscore)
/	~ (tilde)

5. Append the resulting value to your signed URL after `&Signature=`, and return to [To create a signed URL using a canned policy](#) (p. 160) to finish concatenating the parts of your signed URL.

Option 2: To create a signature for an RTMP distribution by using a canned policy (Adobe Flash Player)

1. Use the SHA-1 hash function and RSA to hash and sign the policy statement that you created in the [To create the policy statement for a signed URL that uses a canned policy \(p. 162\)](#) procedure. Use the version of the policy statement that no longer includes whitespace.

For the private key that is required by the hash function, use the private key that is associated with the active trusted signer.

Note

The method that you use to hash and sign the policy statement depends on your programming language and platform. For sample code, see [Code Examples for Creating a Signature for a Signed URL \(p. 189\)](#).

2. Remove whitespace (including tabs and newline characters) from the hashed and signed string.

Continue on to Step 3 if you're using Adobe Flash Player and the stream name is passed in from a web page.

If you're using Adobe Flash Player and if the stream name is not passed in from a web page, skip the rest of this procedure. For example, if you wrote your own player that fetches stream names from within the Adobe Flash .swf file, skip the rest of this procedure.

3. Base64-encode the string using MIME base64 encoding. For more information, see [Section 6.8, Base64 Content-Transfer-Encoding](#) in *RFC 2045, MIME (Multipurpose Internet Mail Extensions) Part One: Format of Internet Message Bodies*.
4. Replace characters that are invalid in a URL query string with characters that are valid. The following table lists invalid and valid characters.

Replace these invalid characters	With these valid characters
+	- (hyphen)
=	_ (underscore)
/	~ (tilde)

5. Some versions of Adobe Flash Player require that you URL-encode the characters ?, =, and &. For information about whether your version of Adobe Flash Player requires this character substitution, refer to the Adobe website.

If your version of Flash does not require URL-encoding those character, skip to Step 6.

If your version of Flash requires URL-encoding those characters, replace them as indicated in the following table. (You already replaced = in the previous step.)

Replace these invalid characters	With this URL encoding
?	%3F
&	%26

6. Append the resulting value to your signed URL after &Signature=, and return to [To create a signed URL using a canned policy \(p. 160\)](#) to finish concatenating the parts of your signed URL.

Creating a Signed URL Using a Custom Policy

Topics

- [Creating a Policy Statement for a Signed URL That Uses a Custom Policy \(p. 168\)](#)
- [Example Policy Statements for a Signed URL That Uses a Custom Policy \(p. 171\)](#)
- [Creating a Signature for a Signed URL That Uses a Custom Policy \(p. 172\)](#)

To create a signed URL using a custom policy, do the following procedure.

To create a signed URL using a custom policy

1. If you're using .NET or Java to create signed URLs, and if you haven't reformatted the private key for your key pair from the default .pem format to a format compatible with .NET or with Java, do so now. For more information, see [Reformatting the CloudFront Private Key \(.NET and Java Only\) \(p. 152\)](#).
2. Concatenate the following values in the specified order, and remove the whitespace (including tabs and newline characters) between the parts. You might have to include escape characters in the string in application code. All values have a type of String. Each part is keyed by number (**1**) to the two examples that follow.

1 *Base URL for the file*

The base URL is the CloudFront URL that you would use to access the file if you were not using signed URLs, including your own query string parameters, if any. For more information about the format of URLs for web distributions, see [Customizing the URL Format for Files in CloudFront \(p. 106\)](#).

The following examples show values that you specify for web distributions.

- The following CloudFront URL is for a file in a web distribution (using the CloudFront domain name). Note that `image.jpg` is in an `images` directory. The path to the file in the URL must match the path to the file on your HTTP server or in your Amazon S3 bucket.

```
http://d1111111abcdef8.cloudfront.net/images/image.jpg
```

- The following CloudFront URL includes a query string:

```
http://d1111111abcdef8.cloudfront.net/images/image.jpg?size=large
```

- The following CloudFront URLs are for files in a web distribution. Both use an alternate domain name; the second one includes a query string:

```
http://www.example.com/images/image.jpg
```

```
http://www.example.com/images/image.jpg?color=red
```

- The following CloudFront URL is for files in a web distribution that uses an alternate domain name and the HTTPS protocol:

```
https://www.example.com/images/image.jpg
```

For RTMP distributions, the following examples are for files in two different video formats, MP4 and FLV:

- **MP4** – `mp4:sydney-vacation.mp4`
- **FLV** – `sydney-vacation`
- **FLV** – `sydney-vacation.flv`

Note

For .flv files, whether you include the .flv filename extension depends on your player. To serve MP3 audio files or H.264/MPEG-4 video files, you might need to prefix the file name with mp3: or mp4:. Some media players can be configured to add the prefix automatically. The media player might also require you to specify the file name without the file extension (for example, sydney-vacation instead of sydney-vacation.mp4).

2 ?

The ? indicates that query string parameters follow the base URL. Include the ? even if you don't have any query string parameters of your own.

3 *Your query string parameters, if any*&

This value is optional. If you want to add your own query string parameters, for example:

color=red&size=medium

then add them after the ? (see 2) and before the Policy parameter. In certain rare circumstances, you might need to put your query string parameters after Key-Pair-Id.

Important

Your parameters cannot be named Policy, Signature, or Key-Pair-Id.

If you add your own parameters, append an & after each one, including the last one.

4 Policy=**base64 encoded version of policy statement**

Your policy statement in JSON format, with white space removed, then base64 encoded. For more information, see [Creating a Policy Statement for a Signed URL That Uses a Custom Policy](#) (p. 168).

The policy statement controls the access that a signed URL grants to a user: the URL of the file (for web distributions) or the stream name (for RTMP distributions), an expiration date and time, an optional date and time that the URL becomes valid, and an optional IP address or range of IP addresses that are allowed to access the file.

5 &Signature=**hashed and signed version of the policy statement**

A hashed, signed, and base64-encoded version of the JSON policy statement. For more information, see [Creating a Signature for a Signed URL That Uses a Custom Policy](#) (p. 172).

6 &Key-Pair-Id=**active CloudFront key pair Id for the key pair that you're using to sign the policy statement**

The ID for an active CloudFront key pair, for example, APKA9ONS7QCOWEXAMPLE. The CloudFront key pair ID tells CloudFront which public key to use to validate the signed URL. CloudFront compares the information in the signature with the information in the policy statement to verify that the URL has not been tampered with.

The key pair ID that you include in CloudFront signed URLs must be the ID of an active key pair for one of your trusted signers:

- **Web distributions** – The key pair must be associated with an AWS account that is one of the trusted signers for the cache behavior.
- **RTMP distributions** – The key pair must be associated with an AWS account that is one of the trusted signers for the distribution.

For more information, see [Specifying the AWS Accounts That Can Create Signed URLs and Signed Cookies \(Trusted Signers\)](#) (p. 150).


```
        "DateGreaterThan":{"AWS:EpochTime":optional beginning date and time in Unix  
time format and UTC},  
        "IpAddress":{"AWS:SourceIp":optional IP address"}  
    }  
  ]  
}
```

Note the following:

- You can include only one statement.
 - Use UTF-8 character encoding.
 - Include all punctuation and parameter names exactly as specified. Abbreviations for parameter names are not accepted.
 - The order of the parameters in the `Condition` section doesn't matter.
 - For information about the values for `Resource`, `DateLessThan`, `DateGreaterThan`, and `IpAddress`, see [Values that You Specify in the Policy Statement for a Signed URL That Uses a Custom Policy](#) (p. 169).
2. Remove all whitespace (including tabs and newline characters) from the policy statement. You might have to include escape characters in the string in application code.
 3. Base64-encode the policy statement using MIME base64 encoding. For more information, see [Section 6.8, Base64 Content-Transfer-Encoding](#) in *RFC 2045, MIME (Multipurpose Internet Mail Extensions) Part One: Format of Internet Message Bodies*.
 4. Replace characters that are invalid in a URL query string with characters that are valid. The following table lists invalid and valid characters.

Replace these invalid characters	With these valid characters
+	- (hyphen)
=	_ (underscore)
/	~ (tilde)

5. Append the resulting value to your signed URL after `Policy=`.
6. Create a signature for the signed URL by hashing, signing, and base64-encoding the policy statement. For more information, see [Creating a Signature for a Signed URL That Uses a Custom Policy](#) (p. 172).

Values that You Specify in the Policy Statement for a Signed URL That Uses a Custom Policy

When you create a policy statement for a custom policy, you specify the following values.

Resource

The value that you specify depends on whether you're creating signed URLs for web or RTMP distributions.

Note

You can specify only one value for `Resource`.

Web distributions (optional but recommended)

The base URL including your query strings, if any, but excluding the `CloudFront Policy`, `Signature`, and `Key-Pair-Id` parameters, for example:

```
http://d111111abcdef8.cloudfront.net/images/horizon.jpg?  
size=large&license=yes
```

Important

If you omit the Resource parameter for a web distribution, users can access all of the files associated with any distribution that is associated with the key pair that you use to create the signed URL.

Note the following:

- **Protocol** – The value must begin with `http://`, `https://`, or `*`.
- **Query string parameters** – If you have no query string parameters, omit the question mark.
- **Wildcard characters** – You can use the wildcard character that matches zero or more characters (`*`) or the wild-card character that matches exactly one character (`?`) anywhere in the string. For example, the value:

```
http://d111111abcdef8.cloudfront.net/*game_download.zip*
```

would include (for example) the following files:

- `http://d111111abcdef8.cloudfront.net/game_download.zip`
- `http://d111111abcdef8.cloudfront.net/example_game_download.zip?license=yes`
- `http://d111111abcdef8.cloudfront.net/test_game_download.zip?license=temp`
- **Alternate domain names** – If you specify an alternate domain name (CNAME) in the URL, you must specify the alternate domain name when referencing the file in your web page or application. Do not specify the Amazon S3 URL for the file.

RTMP distributions

Include only the stream name. For example, if the full URL for a streaming video is:

```
rtmp://s5c39gqb8ow64r.cloudfront.net/videos/cfx/st/mp3_name.mp3
```

then use the following value for Resource:

```
videos/mp3_name
```

Do not include a prefix such as `mp3:` or `mp4:`. Also, depending on the player you're using, you might have to omit the file extension from the value of Resource. For example, you might need to use `sydney-vacation` instead of `sydney-vacation.flv`.

DateLessThan

The expiration date and time for the URL in Unix time format (in seconds) and Coordinated Universal Time (UTC). Do not enclose the value in quotation marks. For information about UTC, see *RFC 3339, Date and Time on the Internet: Timestamps*, <http://tools.ietf.org/html/rfc3339>.

For example, January 1, 2013 10:00 am UTC converts to 1357034400 in Unix time format.

This is the only required parameter in the Condition section. CloudFront requires this value to prevent users from having permanent access to your private content.

For more information, see [When Does CloudFront Check the Expiration Date and Time in a Signed URL?](#) (p. 159)

DateGreaterThan (Optional)

An optional start date and time for the URL in Unix time format (in seconds) and Coordinated Universal Time (UTC). Users are not allowed to access the file before the specified date and time. Do not enclose the value in quotation marks.

IpAddress (Optional)

The IP address of the client making the GET request. Note the following:

- "IPv4 IP address/32"

- ## Important

If you're using a custom policy that includes `IpAddress`, do not enable IPv6 for the distribution. If you want to restrict access to some content by IP address and support IPv6 requests for other content, you can create two distributions. For more information, see [Enable IPv6 \(p. 59\)](#) in the topic [Values That You Specify When You Create or Update a Distribution \(p. 38\)](#).

The following example policy statements show how to control access to a specific file, all of the files in a directory, or all of the files associated with a key pair ID. The examples also show how to control access from an individual IP address or a range of IP addresses, and how to prevent users from using the signed URL after a specified date and time.

For more information, see [Values that You Specify in the Policy Statement for a Signed URL That Uses a Custom Policy \(p. 169\)](#).

- [Example Policy Statement: Accessing One File from a Range of IP Addresses \(p. 171\)](#)
- [Example Policy Statement: Accessing All Files in a Directory from a Range of IP Addresses \(p. 172\)](#)
- [Example Policy Statement: Accessing All Files Associated with a Key Pair ID from One IP Address \(p. 172\)](#)

The following example custom policy in a signed URL specifies that a user can access the file `http://d1111111abcdef8.cloudfront.net/game_download.zip` from IP addresses in the range `192.0.2.0/24` until January 1, 2013 10:00 am UTC:

Example Policy Statement: Accessing All Files in a Directory from a Range of IP Addresses

The following example custom policy allows you to create signed URLs for any file in the `training` directory, as indicated by the `*` wildcard character in the `Resource` parameter. Users can access the file from an IP address in the range `192.0.2.0/24` until January 1, 2013 10:00 am UTC:

```
{
  "Statement": [
    {
      "Resource": "http://d111111abcdef8.cloudfront.net/training/*",
      "Condition": {
        "IpAddress": { "AWS:SourceIp": "192.0.2.0/24" },
        "DateLessThan": { "AWS:EpochTime": 1357034400 }
      }
    }
  ]
}
```

Each signed URL in which you use this policy includes a base URL that identifies a specific file, for example:

`http://d111111abcdef8.cloudfront.net/training/orientation.pdf`

Example Policy Statement: Accessing All Files Associated with a Key Pair ID from One IP Address

The following sample custom policy allows you to create signed URLs for any file associated with any distribution, as indicated by the `*` wildcard character in the `Resource` parameter. The user must use the IP address `192.0.2.10/32`. (The value `192.0.2.10/32` in CIDR notation refers to a single IP address, `192.0.2.10`.) The files are available only from January 1, 2013 10:00 am UTC until January 2, 2013 10:00 am UTC:

```
{
  "Statement": [
    {
      "Resource": "http://*",
      "Condition": {
        "IpAddress": { "AWS:SourceIp": "192.0.2.10/32" },
        "DateGreaterThan": { "AWS:EpochTime": 1357034400 },
        "DateLessThan": { "AWS:EpochTime": 1357120800 }
      }
    }
  ]
}
```

Each signed URL in which you use this policy includes a base URL that identifies a specific file in a specific CloudFront distribution, for example:

`http://d111111abcdef8.cloudfront.net/training/orientation.pdf`

The signed URL also includes a key pair ID, which must be associated with a trusted signer in the distribution (`d111111abcdef8.cloudfront.net`) that you specify in the base URL.

Creating a Signature for a Signed URL That Uses a Custom Policy

The signature for a signed URL that uses a custom policy is a hashed, signed, and base64-encoded version of the policy statement. To create a signature for a custom policy, follow the procedure for your scenario. The version that you choose depends on your distribution type (web or RTMP) and, for RTMP distributions, the media player that you're using (Adobe Flash Player or another media player):

- [Option 1: To create a signature for a web distribution or for an RTMP distribution \(without Adobe Flash Player\) by using a custom policy \(p. 173\)](#)

- [Option 2: To create a signature for an RTMP distribution by using a custom policy \(Adobe Flash Player\) \(p. 173\)](#)

For additional information and examples of how to hash, sign, and encode the policy statement, see:

- [Using a Linux Command and OpenSSL for Base64-Encoding and Encryption \(p. 189\)](#)
- [Code Examples for Creating a Signature for a Signed URL \(p. 189\)](#)

Option 1: To create a signature for a web distribution or for an RTMP distribution (without Adobe Flash Player) by using a custom policy

1. Use the SHA-1 hash function and RSA to hash and sign the JSON policy statement that you created in the procedure [To create the policy statement for a signed URL that uses a custom policy \(p. 168\)](#). Use the version of the policy statement that no longer includes whitespace but that has not yet been base64-encoded.

For the private key that is required by the hash function, use the private key that is associated with the active trusted signer.

Note

The method that you use to hash and sign the policy statement depends on your programming language and platform. For sample code, see [Code Examples for Creating a Signature for a Signed URL \(p. 189\)](#).

2. Remove whitespace (including tabs and newline characters) from the hashed and signed string.
3. Base64-encode the string using MIME base64 encoding. For more information, see [Section 6.8, Base64 Content-Transfer-Encoding](#) in *RFC 2045, MIME (Multipurpose Internet Mail Extensions) Part One: Format of Internet Message Bodies*.
4. Replace characters that are invalid in a URL query string with characters that are valid. The following table lists invalid and valid characters.

Replace these invalid characters	With these valid characters
+	- (hyphen)
=	_ (underscore)
/	~ (tilde)

5. Append the resulting value to your signed URL after `&Signature=`, and return to [To create a signed URL using a custom policy \(p. 166\)](#) to finish concatenating the parts of your signed URL.

Option 2: To create a signature for an RTMP distribution by using a custom policy (Adobe Flash Player)

1. Use the SHA-1 hash function and RSA to hash and sign the JSON policy statement that you created in the [To create the policy statement for a signed URL that uses a custom policy \(p. 168\)](#) procedure. Use the version of the policy statement that no longer includes whitespace but that has not yet been base64-encoded.

For the private key that is required by the hash function, use the private key that is associated with the active trusted signer.

Note

The method that you use to hash and sign the policy statement depends on your programming language and platform. For sample code, see [Code Examples for Creating a Signature for a Signed URL \(p. 189\)](#).

2. Remove whitespace (including tabs and newline characters) from the hashed and signed string.

Continue on to Step 3 if the stream name is passed in from a web page.

If the stream name is not passed in from a web page, skip the rest of this procedure. For example, if you wrote your own player that fetches stream names from within the Adobe Flash .swf file, skip the rest of this procedure.

3. Base64-encode the string using MIME base64 encoding. For more information, see [Section 6.8, Base64 Content-Transfer-Encoding](#) in *RFC 2045, MIME (Multipurpose Internet Mail Extensions) Part One: Format of Internet Message Bodies*.
4. Replace characters that are invalid in a URL query string with characters that are valid. The following table lists invalid and valid characters.

Replace these invalid characters	With these valid characters
+	- (hyphen)
=	_ (underscore)
/	~ (tilde)

5. Some versions of Adobe Flash Player require that you URL-encode the characters ?, =, and &. For information about whether your version of Adobe Flash Player requires this character substitution, refer to the Adobe website.

If your version of Adobe Flash Player does not require that you URL-encode the characters ?, =, and &, skip to Step 6.

If your version of Adobe Flash Player requires URL-encoding those characters, replace them as indicated in the following table. (You already replaced = in the previous step.)

Replace these invalid characters	With this URL encoding
?	%3F
&	%26

6. Append the resulting value to your signed URL after &Signature=, and return to [To create a signed URL using a custom policy \(p. 166\)](#) to finish concatenating the parts of your signed URL.

Using Signed Cookies

CloudFront signed cookies allow you to control who can access your content when you don't want to change your current URLs or when you want to provide access to multiple restricted files, for example, all of the files in the subscribers' area of a website. This topic explains the considerations when using signed cookies and describes how to set signed cookies using canned and custom policies.

Topics

- [Choosing Between Canned and Custom Policies for Signed Cookies \(p. 175\)](#)
- [How Signed Cookies Work \(p. 175\)](#)
- [Preventing Misuse of Signed Cookies \(p. 176\)](#)
- [When Does CloudFront Check the Expiration Date and Time in a Signed Cookie? \(p. 177\)](#)
- [Sample Code and Third-Party Tools \(p. 177\)](#)
- [Setting Signed Cookies Using a Canned Policy \(p. 177\)](#)
- [Setting Signed Cookies Using a Custom Policy \(p. 182\)](#)

Choosing Between Canned and Custom Policies for Signed Cookies

When you create a signed cookie, you write a policy statement in JSON format that specifies the restrictions on the signed cookie, for example, how long the cookie is valid. You can use canned policies or custom policies. The following table compares canned and custom policies:

Description	Canned Policy	Custom Policy
You can reuse the policy statement for multiple files. To reuse the policy statement, you must use wildcard characters in the Resource object. For more information, see Values That You Specify in the Policy Statement for a Custom Policy for Signed Cookies (p. 185) .)	No	Yes
You can specify the date and time that users can begin to access your content	No	Yes (optional)
You can specify the date and time that users can no longer access your content	Yes	Yes
You can specify the IP address or range of IP addresses of the users who can access your content	No	Yes (optional)

For information about creating signed cookies using a canned policy, see [Setting Signed Cookies Using a Canned Policy \(p. 177\)](#).

For information about creating signed cookies using a custom policy, see [Setting Signed Cookies Using a Custom Policy \(p. 182\)](#).

How Signed Cookies Work

Here's an overview of how you configure CloudFront for signed cookies and how CloudFront responds when a user submits a request that contains a signed cookie.

1. In your CloudFront distribution, you specify one or more trusted signers, which are the AWS accounts that you want to have permission to create signed URLs and signed cookies.

For more information, see [Specifying the AWS Accounts That Can Create Signed URLs and Signed Cookies \(Trusted Signers\)](#) (p. 150).

2. You develop your application to determine whether a user should have access to your content and, if so, to send three `Set-Cookie` headers to the viewer. (Each `Set-Cookie` header can contain only one name-value pair, and a CloudFront signed cookie requires three name-value pairs.) You must send the `Set-Cookie` headers to the viewer before the viewer requests your private content. If you set a short expiration time on the cookie, you might also want to send three more `Set-Cookie` headers in response to subsequent requests, so that the user continues to have access.

Typically, your CloudFront distribution will have at least two cache behaviors, one that doesn't require authentication and one that does. The error page for the secure portion of the site includes a redirector or a link to a login page.

If you configure your distribution to cache files based on cookies, CloudFront doesn't cache separate files based on the attributes in signed cookies.

3. A user signs in to your website and either pays for content or meets some other requirement for access.
4. Your application returns the `Set-Cookie` headers in the response, and the viewer stores the name-value pairs.
5. The user requests a file.

The user's browser or other viewer gets the name-value pairs from step 4 and adds them to the request in a `Cookie` header. This is the signed cookie.

6. CloudFront uses the public key to validate the signature in the signed cookie and to confirm that the cookie hasn't been tampered with. If the signature is invalid, the request is rejected.

If the signature in the cookie is valid, CloudFront looks at the policy statement in the cookie (or constructs one if you're using a canned policy) to confirm that the request is still valid. For example, if you specified a beginning and ending date and time for the cookie, CloudFront confirms that the user is trying to access your content during the time period that you want to allow access.

If the request meets the requirements in the policy statement, CloudFront serves your content as it does for content that isn't restricted: it determines whether the file is already in the edge cache, forwards the request to the origin if necessary, and returns the file to the user.

Preventing Misuse of Signed Cookies

If you specify the `Domain` parameter in a `Set-Cookie` header, specify the most precise value possible to reduce the potential for access by someone with the same root domain name. For example, `app.example.com` is preferable to `example.com`, especially when you don't control `example.com`. This helps prevent someone from accessing your content from `www.example.com`.

To help prevent this type of attack, do the following:

- Exclude the `Expires` and `Max-Age` cookie attributes, so that the `Set-Cookie` header creates a session cookie. Session cookies are automatically deleted when the user closes the browser, which reduces the possibility of someone getting unauthorized access to your content.
- Include the `Secure` attribute, so that the cookie is encrypted when a viewer includes it in a request.
- When possible, use a custom policy and include the IP address of the viewer.
- In the `CloudFront-Expires` attribute, specify the shortest reasonable expiration time based on how long you want users to have access to your content.

When Does CloudFront Check the Expiration Date and Time in a Signed Cookie?

To determine whether a signed cookie is still valid, CloudFront checks the expiration date and time in the cookie at the time of the HTTP request. If a client begins to download a large file immediately before the expiration time, the download should complete even if the expiration time passes during the download. If the TCP connection drops and the client tries to restart the download after the expiration time passes, the download will fail.

If a client uses Range GETs to get a file in smaller pieces, any GET request that occurs after the expiration time passes will fail. For more information about Range GETs, see [How CloudFront Processes Partial Requests for an Object \(Range GETs\)](#) (p. 285).

Sample Code and Third-Party Tools

The sample code for private content shows only how to create the signature for signed URLs. However, the process for creating a signature for a signed cookie is very similar, so much of the sample code is still relevant. For more information, see the following topics:

- [Create a URL Signature Using Perl](#) (p. 190)
- [Create a URL Signature Using PHP](#) (p. 198)
- [Create a URL Signature Using C# and the .NET Framework](#) (p. 203)
- [Create a URL Signature Using Java](#) (p. 208)

Setting Signed Cookies Using a Canned Policy

To set a signed cookie by using a canned policy, complete the following steps. To create the signature, see [Creating a Signature for a Signed Cookie That Uses a Canned Policy](#).

To set a signed cookie using a canned policy

1. If you're using .NET or Java to create signed cookies, and if you haven't reformatted the private key for your key pair from the default .pem format to a format compatible with .NET or with Java, do so now. For more information, see [Reformatting the CloudFront Private Key \(.NET and Java Only\)](#) (p. 152).
2. Program your application to send three Set-Cookie headers to approved viewers. You need three Set-Cookie headers because each Set-Cookie header can contain only one name-value pair, and a CloudFront signed cookie requires three name-value pairs. The name-value pairs are: CloudFront-Expires, CloudFront-Signature, and CloudFront-Key-Pair-Id. The values must be present on the viewer before a user makes the first request for an file that you want to control access to.

Note

In general, we recommend that you exclude Expires and Max-Age attributes. Excluding the attributes causes the browser to delete the cookie when the user closes the browser, which reduces the possibility of someone getting unauthorized access to your content. For more information, see [Preventing Misuse of Signed Cookies](#) (p. 176).

The names of cookie attributes are case sensitive.

Line breaks are included only to make the attributes more readable.

```
Set-Cookie:  
CloudFront-Expires=date and time in Unix time format (in seconds) and Coordinated  
Universal Time (UTC);  
Domain=optional domain name;
```

```
Path=/optional directory path;  
Secure;  
HttpOnly  
  
Set-Cookie:  
CloudFront-Signature=hashed and signed version of the policy statement;  
Domain=optional domain name;  
Path=/optional directory path;  
Secure;  
HttpOnly  
  
Set-Cookie:  
CloudFront-Key-Pair-Id=active CloudFront key pair Id for the key pair that you are  
using to generate the signature;  
Domain=optional domain name;  
Path=/optional directory path;  
Secure;  
HttpOnly
```

(Optional) Domain

The domain name for the requested file. If you don't specify a `Domain` attribute, the default value is the domain name in the URL, and it applies only to the specified domain name, not to subdomains. If you specify a `Domain` attribute, it also applies to subdomains. A leading dot in the domain name (for example, `Domain=.example.com`) is optional. In addition, if you specify a `Domain` attribute, the domain name in the URL and the value of the `Domain` attribute must match.

You can specify the domain name that CloudFront assigned to your distribution, for example, `d111111abcdef8.cloudfront.net`, but you can't specify `*.cloudfront.net` for the domain name.

If you want to use an alternate domain name such as `example.com` in URLs, you must add the alternate domain name to your distribution regardless of whether you specify the `Domain` attribute. For more information, see [Alternate Domain Names \(CNAMEs\) \(p. 54\)](#) in the topic [Values That You Specify When You Create or Update a Distribution \(p. 38\)](#).

(Optional) Path

The path for the requested file. If you don't specify a `Path` attribute, the default value is the path in the URL.

Secure

Requires that the viewer encrypt cookies before sending a request. We recommend that you send the `Set-Cookie` header over an HTTPS connection to ensure that the cookie attributes are protected from man-in-the-middle attacks.

HttpOnly

Requires that the viewer send the cookie only in HTTP or HTTPS requests.

CloudFront-Expires

Specify the expiration date and time in Unix time format (in seconds) and Coordinated Universal Time (UTC). For example, January 1, 2013 10:00 am UTC converts to 1357034400 in Unix time format. To use epoch time, use a 32-bit integer for a date which can be no later than 2147483647 (January 19th, 2038 at 03:14:07 UTC). For information about UTC, see *RFC 3339, Date and Time on the Internet: Timestamps*, <http://tools.ietf.org/html/rfc3339>.

CloudFront-Signature

A hashed, signed, and base64-encoded version of a JSON policy statement. For more information, see [Creating a Signature for a Signed Cookie That Uses a Canned Policy \(p. 179\)](#).

CloudFront-Key-Pair-Id

The ID for an active CloudFront key pair, for example, APKA9ONS7QCOWEXAMPLE. The CloudFront key pair ID tells CloudFront which public key to use to validate the signed cookie. CloudFront compares the information in the signature with the information in the policy statement to verify that the URL has not been tampered with.

The key pair ID that you include in CloudFront signed cookies must be associated with an AWS account that is one of the trusted signers for the cache behavior.

For more information, see [Specifying the AWS Accounts That Can Create Signed URLs and Signed Cookies \(Trusted Signers\)](#) (p. 150).

If you make a key pair inactive while rotating CloudFront key pairs, you must update your application to use a new active key pair for one of your trusted signers. For more information about rotating key pairs, see [Rotating CloudFront Key Pairs](#) (p. 155).

The following example shows Set-Cookie headers for one signed cookie when you're using the domain name that is associated with your distribution in the URLs for your files:

```
Set-Cookie: CloudFront-Expires=1426500000; Domain=d111111abcdef8.cloudfront.net; Path=/images/*; Secure; HttpOnly
Set-Cookie: CloudFront-Signature=yXrSigyQoeE4FBI4eMKF6ho-CA8_; Domain=d111111abcdef8.cloudfront.net; Path=/images/*; Secure; HttpOnly
Set-Cookie: CloudFront-Key-Pair-Id=APKA9ONS7QCOWEXAMPLE; Domain=d111111abcdef8.cloudfront.net; Path=/images/*; Secure; HttpOnly
```

The following example shows Set-Cookie headers for one signed cookie when you're using the alternate domain name example.org in the URLs for your files:

```
Set-Cookie: CloudFront-Expires=1426500000; Domain=example.org; Path=/images/*; Secure; HttpOnly
Set-Cookie: CloudFront-Signature=yXrSigyQoeE4FBI4eMKF6ho-CA8_; Domain=example.org; Path=/images/*; Secure; HttpOnly
Set-Cookie: CloudFront-Key-Pair-Id=APKA9ONS7QCOWEXAMPLE; Domain=example.org; Path=/images/*; Secure; HttpOnly
```

If you want to use an alternate domain name such as example.com in URLs, you must add the alternate domain name to your distribution regardless of whether you specify the Domain attribute. For more information, see [Alternate Domain Names \(CNAMEs\)](#) (p. 54) in the topic [Values That You Specify When You Create or Update a Distribution](#) (p. 38).

Creating a Signature for a Signed Cookie That Uses a Canned Policy

To create the signature for a signed cookie that uses a canned policy, do the following:

1. Create a policy statement. See [Creating a Policy Statement for a Signed Cookie That Uses a Canned Policy](#) (p. 179).
2. Sign the policy statement to create a signature. See [Signing the Policy Statement to Create a Signature for a Signed Cookie That Uses a Canned Policy](#) (p. 181).

Creating a Policy Statement for a Signed Cookie That Uses a Canned Policy

When you set a signed cookie that uses a canned policy, the CloudFront-Signature attribute is a hashed and signed version of a policy statement. For signed cookies that use a canned policy, you don't include the policy statement in the Set-Cookie header, as you do for signed cookies that use a custom policy. To create the policy statement, do the following procedure.

To create a policy statement for a signed cookie that uses a canned policy

1. Construct the policy statement using the following JSON format and using UTF-8 character encoding. Include all punctuation and other literal values exactly as specified. For information about the `Resource` and `DateLessThan` parameters, see [Values That You Specify in the Policy Statement for a Canned Policy for Signed Cookies](#) (p. 180).

```
{
  "Statement": [
    {
      "Resource": "base URL or stream name",
      "Condition": {
        "DateLessThan": {
          "AWS:EpochTime": ending date and time in Unix time format and UTC
        }
      }
    }
  ]
}
```

2. Remove all whitespace (including tabs and newline characters) from the policy statement. You might have to include escape characters in the string in application code.

Values That You Specify in the Policy Statement for a Canned Policy for Signed Cookies

When you create a policy statement for a canned policy, you specify the following values:

Resource

The base URL including your query strings, if any, for example:

```
http://d111111abcdef8.cloudfront.net/images/horizon.jpg?
size=large&license=yes
```

You can specify only one value for `Resource`.

Note the following:

- **Protocol** – The value must begin with `http://` or `https://`.
- **Query string parameters** – If you have no query string parameters, omit the question mark.
- **Alternate domain names** – If you specify an alternate domain name (CNAME) in the URL, you must specify the alternate domain name when referencing the file in your web page or application. Do not specify the Amazon S3 URL for the file.

DateLessThan

The expiration date and time for the URL in Unix time format (in seconds) and Coordinated Universal Time (UTC). Do not enclose the value in quotation marks.

For example, March 16, 2015 10:00 am UTC converts to 1426500000 in Unix time format.

This value must match the value of the `CloudFront-Expires` attribute in the `Set-Cookie` header. Do not enclose the value in quotation marks.

For more information, see [When Does CloudFront Check the Expiration Date and Time in a Signed Cookie?](#) (p. 177).

Example Policy Statement for a Canned Policy

When you use the following example policy statement in a signed cookie, a user can access the file `http://d111111abcdef8.cloudfront.net/horizon.jpg` until March 16, 2015 10:00 am UTC:

```
{
  "Statement": [
    {
      "Resource": "http://d111111abcdef8.cloudfront.net/horizon.jpg?
size=large&license=yes",
      "Condition": {
        "DateLessThan": {
          "AWS:EpochTime": 1426500000
        }
      }
    }
  ]
}
```

Signing the Policy Statement to Create a Signature for a Signed Cookie That Uses a Canned Policy

To create the value for the `CloudFront-Signature` attribute in a `Set-Cookie` header, you hash and sign the policy statement that you created in [To create a policy statement for a signed cookie that uses a canned policy \(p. 180\)](#).

For additional information and examples of how to hash, sign, and encode the policy statement, see the following topics:

- [Using a Linux Command and OpenSSL for Base64-Encoding and Encryption \(p. 189\)](#)
- [Code Examples for Creating a Signature for a Signed URL \(p. 189\)](#)

To create a signature for a signed cookie using a canned policy

1. Use the SHA-1 hash function and RSA to hash and sign the policy statement that you created in the procedure [To create a policy statement for a signed cookie that uses a canned policy \(p. 180\)](#). Use the version of the policy statement that no longer includes whitespace.

For the private key that is required by the hash function, use the private key that is associated with the active trusted signer.

Note

The method that you use to hash and sign the policy statement depends on your programming language and platform. For sample code, see [Code Examples for Creating a Signature for a Signed URL \(p. 189\)](#).

2. Remove whitespace (including tabs and newline characters) from the hashed and signed string.
3. Base64-encode the string using MIME base64 encoding. For more information, see [Section 6.8, Base64 Content-Transfer-Encoding](#) in *RFC 2045, MIME (Multipurpose Internet Mail Extensions) Part One: Format of Internet Message Bodies*.
4. Replace characters that are invalid in a URL query string with characters that are valid. The following table lists invalid and valid characters.

Replace these invalid characters	With these valid characters
+	- (hyphen)
=	_ (underscore)
/	~ (tilde)

5. Include the resulting value in the `Set-Cookie` header for the `CloudFront-Signature` name-value pair. Then return to [To set a signed cookie using a canned policy \(p. 177\)](#) add the `Set-Cookie` header for `CloudFront-Key-Pair-Id`.

Setting Signed Cookies Using a Custom Policy

Topics

- [Creating a Policy Statement for a Signed Cookie That Uses a Custom Policy \(p. 184\)](#)
- [Example Policy Statements for a Signed Cookie That Uses a Custom Policy \(p. 186\)](#)
- [Creating a Signature for a Signed Cookie That Uses a Custom Policy \(p. 188\)](#)

To set a signed cookie that uses a custom policy, do the following procedure.

To set a signed cookie using a custom policy

1. If you're using .NET or Java to create signed URLs, and if you haven't reformatted the private key for your key pair from the default .pem format to a format compatible with .NET or with Java, do so now. For more information, see [Reformatting the CloudFront Private Key \(.NET and Java Only\) \(p. 152\)](#).
2. Program your application to send three Set-Cookie headers to approved viewers. You need three Set-Cookie headers because each Set-Cookie header can contain only one name-value pair, and a CloudFront signed cookie requires three name-value pairs. The name-value pairs are: CloudFront-Policy, CloudFront-Signature, and CloudFront-Key-Pair-Id. The values must be present on the viewer before a user makes the first request for a file that you want to control access to.

Note

In general, we recommend that you exclude Expires and Max-Age attributes. This causes the browser to delete the cookie when the user closes the browser, which reduces the possibility of someone getting unauthorized access to your content. For more information, see [Preventing Misuse of Signed Cookies \(p. 176\)](#).

The names of cookie attributes are case sensitive.

Line breaks are included only to make the attributes more readable.

```
Set-Cookie:  
CloudFront-Policy=base64 encoded version of the policy statement;  
Domain=optional domain name;  
Path=/optional directory path;  
Secure;  
HttpOnly  
  
Set-Cookie:  
CloudFront-Signature=hashed and signed version of the policy statement;  
Domain=optional domain name;  
Path=/optional directory path;  
Secure;  
HttpOnly  
  
Set-Cookie:  
CloudFront-Key-Pair-Id=active CloudFront key pair Id for the key pair that you are  
using to generate the signature;  
Domain=optional domain name;  
Path=/optional directory path;  
Secure;  
HttpOnly
```

(Optional) Domain

The domain name for the requested file. If you don't specify a `Domain` attribute, the default value is the domain name in the URL, and it applies only to the specified domain name, not to subdomains. If you specify a `Domain` attribute, it also applies to subdomains. A leading dot in the domain name (for example, `Domain=.example.com`) is optional. In addition, if you specify a `Domain` attribute, the domain name in the URL and the value of the `Domain` attribute must match.

You can specify the domain name that CloudFront assigned to your distribution, for example, `d111111abcdef8.cloudfront.net`, but you can't specify `*.cloudfront.net` for the domain name.

If you want to use an alternate domain name such as `example.com` in URLs, you must add the alternate domain name to your distribution regardless of whether you specify the `Domain` attribute. For more information, see [Alternate Domain Names \(CNAMES\) \(p. 54\)](#) in the topic [Values That You Specify When You Create or Update a Distribution \(p. 38\)](#).

(Optional) Path

The path for the requested file. If you don't specify a `Path` attribute, the default value is the path in the URL.

Secure

Requires that the viewer encrypt cookies before sending a request. We recommend that you send the `Set-Cookie` header over an HTTPS connection to ensure that the cookie attributes are protected from man-in-the-middle attacks.

HttpOnly

Requires that the viewer send the cookie only in HTTP or HTTPS requests.

CloudFront-Policy

Your policy statement in JSON format, with white space removed, then base64 encoded. For more information, see [Creating a Signature for a Signed Cookie That Uses a Custom Policy \(p. 188\)](#).

The policy statement controls the access that a signed cookie grants to a user: the files that the user can access, an expiration date and time, an optional date and time that the URL becomes valid, and an optional IP address or range of IP addresses that are allowed to access the file.

CloudFront-Signature

A hashed, signed, and base64-encoded version of the JSON policy statement. For more information, see [Creating a Signature for a Signed Cookie That Uses a Custom Policy \(p. 188\)](#).

CloudFront-Key-Pair-Id

The ID for an active CloudFront key pair, for example, `APKA9ONS7QCOWEXAMPLE`. The CloudFront key pair ID tells CloudFront which public key to use to validate the signed cookie. CloudFront compares the information in the signature with the information in the policy statement to verify that the URL has not been tampered with.

The key pair ID that you include in CloudFront signed cookies must be associated with an AWS account that is one of the trusted signers for the cache behavior.

For more information, see [Specifying the AWS Accounts That Can Create Signed URLs and Signed Cookies \(Trusted Signers\) \(p. 150\)](#).

If you make a key pair inactive while rotating CloudFront key pairs, you must update your application to use a new active key pair for one of your trusted signers. For more information about rotating key pairs, see [Rotating CloudFront Key Pairs \(p. 155\)](#).

Example Set-Cookie headers for one signed cookie when you're using the domain name that is associated with your distribution in the URLs for your files:

```
Set-Cookie: CloudFront-  
Policy=eyJTdGF0ZWlbnQlOlt7IlJlc291cmNlIjoiaHR0cDovL2QxMTEzMTEhYmNkZWY4LmNsb3VkZnJvbnQubmV0L2dhbWVfZG93  
Domain=d111111abcdef8.cloudfront.net; Path=/; Secure; HttpOnly  
Set-Cookie: CloudFront-Signature=dtKhpJ3aUYxqDIwepczPiDb9NXQ_  
Domain=d111111abcdef8.cloudfront.net; Path=/; Secure; HttpOnly  
Set-Cookie: CloudFront-Key-Pair-Id=APKA9ONS7QCOWEXAMPLE;  
Domain=d111111abcdef8.cloudfront.net; Path=/; Secure; HttpOnly
```

Example Set-Cookie headers for one signed cookie when you're using the alternate domain name example.org in the URLs for your files:

```
Set-Cookie: CloudFront-  
Policy=eyJTdGF0ZWlbnQlOlt7IlJlc291cmNlIjoiaHR0cDovL2QxMTEzMTEhYmNkZWY4LmNsb3VkZnJvbnQubmV0L2dhbWVfZG93  
Domain=example.org; Path=/; Secure; HttpOnly  
Set-Cookie: CloudFront-Signature=dtKhpJ3aUYxqDIwepczPiDb9NXQ_; Domain=example.org; Path=/;  
Secure; HttpOnly  
Set-Cookie: CloudFront-Key-Pair-Id=APKA9ONS7QCOWEXAMPLE; Domain=example.org; Path=/  
Secure; HttpOnly
```

If you want to use an alternate domain name such as example.com in URLs, you must add the alternate domain name to your distribution regardless of whether you specify the Domain attribute. For more information, see [Alternate Domain Names \(CNAMEs\) \(p. 54\)](#) in the topic [Values That You Specify When You Create or Update a Distribution \(p. 38\)](#).

Creating a Policy Statement for a Signed Cookie That Uses a Custom Policy

To create a policy statement for a custom policy, do the following procedure. For several example policy statements that control access to files in a variety of ways, see [Example Policy Statements for a Signed Cookie That Uses a Custom Policy \(p. 186\)](#).

To create the policy statement for a signed cookie that uses a custom policy

1. Construct the policy statement using the following JSON format.

```
{  
  "Statement": [  
    {  
      "Resource": "URL of the file",  
      "Condition": {  
        "DateLessThan": { "AWS:EpochTime": required ending date and time in Unix time  
format and UTC },  
        "DateGreaterThan": { "AWS:EpochTime": optional beginning date and time in Unix  
time format and UTC },  
        "IpAddress": { "AWS:SourceIp": optional IP address }  
      }  
    }  
  ]  
}
```

Note the following:

- You can include only one statement.
- Use UTF-8 character encoding.
- Include all punctuation and parameter names exactly as specified. Abbreviations for parameter names are not accepted.
- The order of the parameters in the Condition section doesn't matter.

- For information about the values for `Resource`, `DateLessThan`, `DateGreaterThan`, and `IpAddress`, see [Values That You Specify in the Policy Statement for a Custom Policy for Signed Cookies](#) (p. 185).
2. Remove all whitespace (including tabs and newline characters) from the policy statement. You might have to include escape characters in the string in application code.
 3. Base64-encode the policy statement using MIME base64 encoding. For more information, see [Section 6.8, Base64 Content-Transfer-Encoding](#) in *RFC 2045, MIME (Multipurpose Internet Mail Extensions) Part One: Format of Internet Message Bodies*.
 4. Replace characters that are invalid in a URL query string with characters that are valid. The following table lists invalid and valid characters.

Replace these invalid characters	With these valid characters
+	- (hyphen)
=	_ (underscore)
/	~ (tilde)

5. Include the resulting value in your `Set-Cookie` header after `CloudFront-Policy=`.
6. Create a signature for the `Set-Cookie` header for `CloudFront-Signature` by hashing, signing, and base64-encoding the policy statement. For more information, see [Creating a Signature for a Signed Cookie That Uses a Custom Policy](#) (p. 188).

Values That You Specify in the Policy Statement for a Custom Policy for Signed Cookies

When you create a policy statement for a custom policy, you specify the following values.

Resource

The base URL including your query strings, if any:

```
http://d111111abcdef8.cloudfront.net/images/horizon.jpg?  
size=large&license=yes
```

Important

If you omit the `Resource` parameter, users can access all of the files associated with any distribution that is associated with the key pair that you use to create the signed URL.

You can specify only one value for `Resource`.

Note the following:

- **Protocol** – The value must begin with `http://` or `https://`.
- **Query string parameters** – If you have no query string parameters, omit the question mark.
- **Wildcards** – You can use the wildcard character that matches zero or more characters (*) or the wild-card character that matches exactly one character (?) anywhere in the string. For example, the value:

```
http://d111111abcdef8.cloudfront.net/*game_download.zip*
```

would include (for example) the following files:

- `http://d111111abcdef8.cloudfront.net/game_download.zip`
- `http://d111111abcdef8.cloudfront.net/example_game_download.zip?license=yes`
- `http://d111111abcdef8.cloudfront.net/test_game_download.zip?license=temp`

- **Alternate domain names** – If you specify an alternate domain name (CNAME) in the URL, you must specify the alternate domain name when referencing the file in your web page or application. Do not specify the Amazon S3 URL for the file.

DateLessThan

The expiration date and time for the URL in Unix time format (in seconds) and Coordinated Universal Time (UTC). Do not enclose the value in quotation marks.

For example, March 16, 2015 10:00 am UTC converts to 1426500000 in Unix time format.

For more information, see [When Does CloudFront Check the Expiration Date and Time in a Signed Cookie?](#) (p. 177).

DateGreaterThan (Optional)

An optional start date and time for the URL in Unix time format (in seconds) and Coordinated Universal Time (UTC). Users are not allowed to access the file before the specified date and time. Do not enclose the value in quotation marks.

IpAddress (Optional)

The IP address of the client making the GET request. Note the following:

- To allow any IP address to access the file, omit the `IpAddress` parameter.
- You can specify either one IP address or one IP address range. For example, you can't set the policy to allow access if the client's IP address is in one of two separate ranges.
- To allow access from a single IP address, you specify:

`"IPv4 IP address/32"`

- You must specify IP address ranges in standard IPv4 CIDR format (for example, `192.0.2.0/24`). For more information, go to *RFC 4632, Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan*, <http://tools.ietf.org/html/rfc4632>.

Important

IP addresses in IPv6 format, such as `2001:0db8:85a3:0000:0000:8a2e:0370:7334`, are not supported.

If you're using a custom policy that includes `IpAddress`, do not enable IPv6 for the distribution. If you want to restrict access to some content by IP address and support IPv6 requests for other content, you can create two distributions. For more information, see [Enable IPv6](#) (p. 59) in the topic [Values That You Specify When You Create or Update a Distribution](#) (p. 38).

Example Policy Statements for a Signed Cookie That Uses a Custom Policy

The following example policy statements show how to control access to a specific file, all of the files in a directory, or all of the files associated with a key pair ID. The examples also show how to control access from an individual IP address or a range of IP addresses, and how to prevent users from using the signed cookie after a specified date and time.

If you copy and paste any of these examples, remove any whitespace (including tabs and newline characters), replace the values with your own values, and include a newline character after the closing brace `}`.

For more information, see [Values That You Specify in the Policy Statement for a Custom Policy for Signed Cookies](#) (p. 185).

Topics

- [Example Policy Statement: Accessing One File from a Range of IP Addresses](#) (p. 187)
- [Example Policy Statement: Accessing All Files in a Directory from a Range of IP Addresses](#) (p. 187)

- [Example Policy Statement: Accessing All Files Associated with a Key Pair ID from One IP Address \(p. 187\)](#)

Example Policy Statement: Accessing One File from a Range of IP Addresses

The following example custom policy in a signed cookie specifies that a user can access the file `http://d111111abcdef8.cloudfront.net/game_download.zip` from IP addresses in the range `192.0.2.0/24` until January 1, 2013 10:00 am UTC:

```
{
  "Statement": [
    {
      "Resource": "http://d111111abcdef8.cloudfront.net/game_download.zip",
      "Condition": {
        "IpAddress": { "AWS:SourceIp": "192.0.2.0/24" },
        "DateLessThan": { "AWS:EpochTime": 1357034400 }
      }
    }
  ]
}
```

Example Policy Statement: Accessing All Files in a Directory from a Range of IP Addresses

The following example custom policy allows you to create signed cookies for any file in the `training` directory, as indicated by the `*` wildcard character in the `Resource` parameter. Users can access the file from an IP address in the range `192.0.2.0/24` until January 1, 2013 10:00 am UTC:

```
{
  "Statement": [
    {
      "Resource": "http://d111111abcdef8.cloudfront.net/training/*",
      "Condition": {
        "IpAddress": { "AWS:SourceIp": "192.0.2.0/24" },
        "DateLessThan": { "AWS:EpochTime": 1357034400 }
      }
    }
  ]
}
```

Each signed cookie in which you use this policy includes a base URL that identifies a specific file, for example:

`http://d111111abcdef8.cloudfront.net/training/orientation.pdf`

Example Policy Statement: Accessing All Files Associated with a Key Pair ID from One IP Address

The following sample custom policy allows you to set signed cookies for any file associated with any distribution, as indicated by the `*` wildcard character in the `Resource` parameter. The user must use the IP address `192.0.2.10/32`. (The value `192.0.2.10/32` in CIDR notation refers to a single IP address, `192.0.2.10`.) The files are available only from January 1, 2013 10:00 am UTC until January 2, 2013 10:00 am UTC:

```
{
  "Statement": [
    {
      "Resource": "http://*",
      "Condition": {
        "IpAddress": { "AWS:SourceIp": "192.0.2.10/32" },
        "DateGreaterThan": { "AWS:EpochTime": 1357034400 },

```

```
        "DateLessThan":{"AWS:EpochTime":1357120800}  
      }  
    }  
  ]  
}
```

Each signed cookie in which you use this policy includes a base URL that identifies a specific file in a specific CloudFront distribution, for example:

`http://d111111abcdef8.cloudfront.net/training/orientation.pdf`

The signed cookie also includes a key pair ID, which must be associated with a trusted signer in the distribution (d111111abcdef8.cloudfront.net) that you specify in the base URL.

Creating a Signature for a Signed Cookie That Uses a Custom Policy

The signature for a signed cookie that uses a custom policy is a hashed, signed, and base64-encoded version of the policy statement.

For additional information and examples of how to hash, sign, and encode the policy statement, see:

- [Using a Linux Command and OpenSSL for Base64-Encoding and Encryption \(p. 189\)](#)
- [Code Examples for Creating a Signature for a Signed URL \(p. 189\)](#)

To create a signature for a signed cookie by using a custom policy

1. Use the SHA-1 hash function and RSA to hash and sign the JSON policy statement that you created in the procedure [To create the policy statement for a signed URL that uses a custom policy \(p. 168\)](#). Use the version of the policy statement that no longer includes whitespace but that has not yet been base64-encoded.

For the private key that is required by the hash function, use the private key that is associated with the active trusted signer.

Note

The method that you use to hash and sign the policy statement depends on your programming language and platform. For sample code, see [Code Examples for Creating a Signature for a Signed URL \(p. 189\)](#).

2. Remove whitespace (including tabs and newline characters) from the hashed and signed string.
3. Base64-encode the string using MIME base64 encoding. For more information, see [Section 6.8, Base64 Content-Transfer-Encoding](#) in *RFC 2045, MIME (Multipurpose Internet Mail Extensions) Part One: Format of Internet Message Bodies*.
4. Replace characters that are invalid in a URL query string with characters that are valid. The following table lists invalid and valid characters.

Replace these invalid characters	With these valid characters
+	- (hyphen)
=	_ (underscore)
/	~ (tilde)

5. Include the resulting value in the Set-Cookie header for the CloudFront-Signature= name-value pair, and return to [To set a signed cookie using a custom policy \(p. 182\)](#) to add the Set-Cookie header for CloudFront-Key-Pair-Id.

Using a Linux Command and OpenSSL for Base64-Encoding and Encryption

You can use the following Linux command-line command and OpenSSL to hash and sign the policy statement, base64-encode the signature, and replace characters that are not valid in URL query string parameters with characters that are valid.

For information about OpenSSL, go to <http://www.openssl.org>.

```
1 cat policy | 2 tr -d "\n" | tr -d " \t\n\r" | 3 openssl sha1 -sign  
private-key.pem | 4 openssl base64 | 5 tr -- '+=/' '-_~'
```

where:

1 `cat` reads the `policy` file.

2 `tr -d "\n" | tr -d " \t\n\r"` removes the white spaces and newline character that were added by `cat`.

3 OpenSSL hashes the file using SHA-1 and signs it using RSA and the private key file `private-key.pem`.

4 OpenSSL base64-encodes the hashed and signed policy statement.

5 `tr` replaces characters that are not valid in URL query string parameters with characters that are valid.

For code examples that demonstrate creating a signature in several programming languages see [Code Examples for Creating a Signature for a Signed URL](#) (p. 189).

Code Examples for Creating a Signature for a Signed URL

This section includes downloadable application examples that demonstrate how to create signatures for signed URLs. Examples are available in Perl, PHP, C#, and Java. You can use any of the examples to create signed URLs. The Perl script runs on Linux/Mac platforms. The PHP example will work on any server that runs PHP. The C# example uses the .NET Framework.

This post on a third-party blog shows an example of CloudFront signed cookies in Ruby: [Using Cloudfront Signed Cookies](#).

This post on the AWS Developer Blog shows an example of CloudFront signed URLs in Node.js, using a third-party npm module called `aws-cloudfront-sign`: [Creating Amazon CloudFront Signed URLs in Node.js](#)

You can find example code for signed URLs and for signed cookies in a variety of programming languages. Do an internet search on `sample app language cloudfront signed URLs` or on `sample app language cloudfront signed cookies`.

Topics

- [Create a URL Signature Using Perl](#) (p. 190)
- [Create a URL Signature Using PHP](#) (p. 198)

=head1 SYNOPSIS

View the script source for details as to which CPAN packages are required beforehand.

```
cfsign.pl --help
```

```
cfsign.pl --action encode --url http://images.my-website.com/gallery1.zip --policy
sample_policy.json --private-key privkey.pem --key-pair-id mykey
```

Stream signing example:

URL decode example:

To generate an RSA key pair, you can use openssl and the following commands:

```
=head1  OPTIONS
```

```
=item B<--help>
```

```
=item B<--action> [action]
```

```
encode - Generate a signed URL (using a canned policy or a user policy)
decode - Decode a signed URL
```

The URL to en/decode

The stream to en/decode

The path to your private key.

```

The AWS Portal assigned key pair identifier.

=item B<--policy>

The CloudFront policy document.

=item B<--expires>

The Unix epoch time when the URL is to expire. If both this option and
the --policy option are specified, --policy will be used. Otherwise, this
option alone will use a canned policy.

=back

=cut

use strict;
use warnings;

# you might need to use CPAN to get these modules.
# run perl -MCPAN -e "install <module>" to get them.
# The openssl command line will also need to be in your $PATH.
use File::Temp qw/tempfile/;
use File::Slurp;
use Getopt::Long;
use IPC::Open2;
use MIME::Base64 qw(encode_base64 decode_base64);
use Pod::Usage;
use URI;

my $CANNED_POLICY
    = '{"Statement":[{"Resource":"<RESOURCE>","Condition":{"DateLessThan":
{"AWS:EpochTime":<EXPIRES>}}}]}'

my $POLICY_PARAM      = "Policy";
my $EXPIRES_PARAM     = "Expires";
my $SIGNATURE_PARAM   = "Signature";
my $KEY_PAIR_ID_PARAM = "Key-Pair-Id";

my $verbose = 0;
my $policy_filename = "";
my $expires_epoch = 0;
my $action = "";
my $help = 0;
my $key_pair_id = "";
my $url = "";
my $stream = "";
my $private_key_filename = "";

my $result = GetOptions("action=s"      => \$action,
                        "policy=s"      => \$policy_filename,
                        "expires=i"     => \$expires_epoch,
                        "private-key=s" => \$private_key_filename,
                        "key-pair-id=s" => \$key_pair_id,
                        "verbose"       => \$verbose,
                        "help"          => \$help,
                        "url=s"         => \$url,
                        "stream=s"      => \$stream,
                        );

if ($help or !$result) {
    pod2usage(1);
    exit;
}

if ($url eq "" and $stream eq "") {

```

```

    print STDERR "Must include a stream or a URL to encode or decode with the --stream or
--url option\n";
    exit;
}

if ($url ne "" and $stream ne "") {
    print STDERR "Only one of --url and --stream may be specified\n";
    exit;
}

if ($url ne "" and !is_url_valid($url)) {
    exit;
}

if ($stream ne "") {
    exit unless is_stream_valid($stream);

    # The signing mechanism is identical, so from here on just pretend we're
    # dealing with a URL
    $url = $stream;
}

if ($action eq "encode") {
    # The encode action will generate a private content URL given a base URL,
    # a policy file (or an expires timestamp) and a key pair id parameter
    my $private_key;
    my $public_key;
    my $public_key_file;

    my $policy;
    if ($policy_filename eq "") {
        if ($expires_epoch == 0) {
            print STDERR "Must include policy filename with --policy argument or an
expires" .
                                "time using --expires\n";
        }

        $policy = $CANNED_POLICY;
        $policy =~ s/<EXPIRES>/$expires_epoch/g;
        $policy =~ s/<RESOURCE>/$url/g;
    } else {
        if (! -e $policy_filename) {
            print STDERR "Policy file $policy_filename does not exist\n";
            exit;
        }
        $expires_epoch = 0; # ignore if set
        $policy = read_file($policy_filename);
    }

    if ($private_key_filename eq "") {
        print STDERR "You must specific the path to your private key file with --private-
key\n";
        exit;
    }

    if (! -e $private_key_filename) {
        print STDERR "Private key file $private_key_filename does not exist\n";
        exit;
    }

    if ($key_pair_id eq "") {
        print STDERR "You must specify an AWS portal key pair id with --key-pair-id\n";
        exit;
    }

    my $encoded_policy = url_safe_base64_encode($policy);

```

```

my $signature = rsa_shal_sign($policy, $private_key_filename);
my $encoded_signature = url_safe_base64_encode($signature);

my $generated_url = create_url($url, $encoded_policy, $encoded_signature, $key_pair_id,
$expires_epoch);

if ($stream ne "") {
    print "Encoded stream (for use within a swf):\n" . $generated_url . "\n";
    print "Encoded and escaped stream (for use on a webpage):\n" .
escape_url_for_webpage($generated_url) . "\n";
} else {
    print "Encoded URL:\n" . $generated_url . "\n";
}
} elsif ($action eq "decode") {
my $decoded = decode_url($url);
if (!$decoded) {
    print STDERR "Improperly formed URL\n";
    exit;
}

print_decoded_url($decoded);
} else {
    # No action specified, print help. But only if this is run as a program (caller will
    be empty)
    pod2usage(1) unless caller();
}

# Decode a private content URL into its component parts
sub decode_url {
    my $url = shift;

    if ($url =~ /(.*?)\?(.*)/) {
        my $base_url = $1;
        my $params = $2;

        my @unparsed_params = split(/&/, $params);
        my %params = ();
        foreach my $param (@unparsed_params) {
            my ($key, $val) = split(/=/, $param);
            $params{$key} = $val;
        }

        my $encoded_signature = "";
        if (exists $params{$SIGNATURE_PARAM}) {
            $encoded_signature = $params{"Signature"};
        } else {
            print STDERR "Missing Signature URL parameter\n";
            return 0;
        }

        my $encoded_policy = "";
        if (exists $params{$POLICY_PARAM}) {
            $encoded_policy = $params{$POLICY_PARAM};
        } else {
            if (!exists $params{$EXPIRES_PARAM}) {
                print STDERR "Either the Policy or Expires URL parameter needs to be
specified\n";
                return 0;
            }

            my $expires = $params{$EXPIRES_PARAM};

            my $policy = $CANNED_POLICY;
            $policy =~ s/<EXPIRES>/$expires/g;

```

```

my $url_without_cf_params = $url;
$url_without_cf_params =~ s/$SIGNATURE_PARAM=[^&]*&?//g;
$url_without_cf_params =~ s/$POLICY_PARAM=[^&]*&?//g;
$url_without_cf_params =~ s/$EXPIRES_PARAM=[^&]*&?//g;
$url_without_cf_params =~ s/$KEY_PAIR_ID_PARAM=[^&]*&?//g;

if ($url_without_cf_params =~ /(.*?)\?$/) {
    $url_without_cf_params = $1;
}

$policy =~ s/<RESOURCE>/$url_without_cf_params/g;

$encoded_policy = url_safe_base64_encode($policy);
}

my $key = "";
if (exists $params{$KEY_PAIR_ID_PARAM}) {
    $key = $params{$KEY_PAIR_ID_PARAM};
} else {
    print STDERR "Missing $KEY_PAIR_ID_PARAM parameter\n";
    return 0;
}

my $policy = url_safe_base64_decode($encoded_policy);

my %ret = ();
$ret{"base_url"} = $base_url;
$ret{"policy"} = $policy;
$ret{"key"} = $key;

return \%ret;
} else {
    return 0;
}
}

# Print a decoded URL out
sub print_decoded_url {
    my $decoded = shift;

    print "Base URL: \n" . $decoded->{"base_url"} . "\n";
    print "Policy: \n" . $decoded->{"policy"} . "\n";
    print "Key: \n" . $decoded->{"key"} . "\n";
}

# Encode a string with base 64 encoding and replace some invalid URL characters
sub url_safe_base64_encode {
    my ($value) = @_;

    my $result = encode_base64($value);
    $result =~ tr|+|=/_~|;

    return $result;
}

# Decode a string with base 64 encoding. URL-decode the string first
# followed by reversing any special character ("+=/") translation.
sub url_safe_base64_decode {
    my ($value) = @_;

    $value =~ s/%([0-9A-Fa-f]{2})/chr(hex($1))/eg;
    $value =~ tr|_~|+|=/_|;

    my $result = decode_base64($value);

    return $result;
}

```

```

}

# Create a private content URL
sub create_url {
    my ($path, $policy, $signature, $key_pair_id, $expires) = @_;

    my $result;
    my $separator = $path =~ /\?/ ? '&' : '?';
    if ($expires) {
        $result = "$path$separator$EXPIRES_PARAM=$expires&$SIGNATURE_PARAM=$signature&
$KEY_PAIR_ID_PARAM=$key_pair_id";
    } else {
        $result = "$path$separator$POLICY_PARAM=$policy&$SIGNATURE_PARAM=$signature&
$KEY_PAIR_ID_PARAM=$key_pair_id";
    }
    $result =~ s/\n/g;

    return $result;
}

# Sign a document with given private key file.
# The first argument is the document to sign
# The second argument is the name of the private key file
sub rsa_shal_sign {
    my ($to_sign, $pvkFile) = @_;
    print "openssl shal -sign $pvkFile $to_sign\n";

    return write_to_program($pvkFile, $to_sign);
}

# Helper function to write data to a program
sub write_to_program {
    my ($keyfile, $data) = @_;
    unlink "temp_policy.dat" if (-e "temp_policy.dat");
    unlink "temp_sign.dat" if (-e "temp_sign.dat");

    write_file("temp_policy.dat", $data);

    system("openssl dgst -shal -sign \"\$keyfile\" -out temp_sign.dat temp_policy.dat");

    my $output = read_file("temp_sign.dat");

    return $output;
}

# Read a file into a string and return the string
sub read_file {
    my ($file) = @_;

    open(INFILE, "<$file") or die("Failed to open $file: $!");
    my $str = join('', <INFILE>);
    close INFILE;

    return $str;
}

sub is_url_valid {
    my ($url) = @_;

    # HTTP distributions start with http[s]:// and are the correct thing to sign
    if ($url =~ /^https?:\\/\\/){
        return 1;
    } else {
        print STDERR "CloudFront requires absolute URLs for HTTP distributions\n";
        return 0;
    }
}

```

```
}

sub is_stream_valid {
    my ($stream) = @_;

    if ($stream =~ /^rtmp:\// or $stream =~ /^\/?cfx\/st/) {
        print STDERR "Streaming distributions require that only the stream name is signed.
\n";
        print STDERR "The stream name is everything after, but not including, cfx/st/\n";
        return 0;
    } else {
        return 1;
    }
}

# flash requires that the query parameters in the stream name are url
# encoded when passed in through javascript, etc. This sub handles the minimal
# required url encoding.
sub escape_url_for_webpage {
    my ($url) = @_;

    $url =~ s/\?/%3F/g;
    $url =~ s/=/%3D/g;
    $url =~ s/&/%26/g;

    return $url;
}

1;
```

Create a URL Signature Using PHP

Any web server that runs PHP can use this PHP example code to create policy statements and signatures for private CloudFront distributions. The full example creates a functioning web page with signed URL links that play a video stream using CloudFront streaming. You can download the full example at <https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/samples/demo-php.zip>.

You can also create signed URLs by using the `UrlSigner` class in the AWS SDK for PHP. For more information, see [Class UrlSigner](#) in the *AWS SDK for PHP API Reference*.

Note

Creating a URL signature is just one part of the process of serving private content using a signed URL. For more information about the entire process, see [Using Signed URLs \(p. 157\)](#).

Topics

- [Sample: RSA SHA-1 Signature \(p. 198\)](#)
- [Sample: Create Canned Policy \(p. 199\)](#)
- [Sample: Create a Custom Policy \(p. 199\)](#)
- [Full Code Example \(p. 200\)](#)

Sample: RSA SHA-1 Signature

In the following code sample, the function `rsa_sha1_sign` hashes and signs the policy statement. The arguments required are a policy statement and the private key for an AWS account that is a trusted signer for your distribution. Next, the `url_safe_base64_encode` function creates a URL-safe version of the signature.

```
function rsa_sha1_sign($policy, $private_key_filename) {
```

```
$signature = "";

// load the private key
$fp = fopen($private_key_filename, "r");
$priv_key = fread($fp, 8192);
fclose($fp);
$pkeyid = openssl_get_privatekey($priv_key);

// compute signature
openssl_sign($policy, $signature, $pkeyid);

// free the key from memory
openssl_free_key($pkeyid);

return $signature;
}

function url_safe_base64_encode($value) {
    $encoded = base64_encode($value);
    // replace unsafe characters +, = and / with
    // the safe characters -, _ and ~
    return str_replace(
        array('+', '=', '/'),
        array('-', '_', '~'),
        $encoded);
}
```

Sample: Create Canned Policy

The following sample code constructs a *canned* policy statement for the signature. For more information about canned policies, see [Creating a Signed URL Using a Canned Policy \(p. 160\)](#).

Note

The `$expires` variable is a date/time stamp that must be an integer, not a string.

```
function get_canned_policy_stream_name($video_path, $private_key_filename, $key_pair_id,
    $expires) {
    // this policy is well known by CloudFront, but you still need to sign it,
    // since it contains your parameters
    $canned_policy = '{"Statement":[{"Resource":"' . $video_path . '","Condition":
{"DateLessThan":{"AWS:EpochTime":"' . $expires . '}}}]}';

    // sign the canned policy
    $signature = rsa_sha1_sign($canned_policy, $private_key_filename);
    // make the signature safe to be included in a url
    $encoded_signature = url_safe_base64_encode($signature);

    // combine the above into a stream name
    $stream_name = create_stream_name($video_path, null, $encoded_signature, $key_pair_id,
    $expires);
    // url-encode the query string characters to work around a flash player bug
    return encode_query_params($stream_name);
}
```

Sample: Create a Custom Policy

The following sample code constructs a *custom* policy statement for the signature. For more information about custom policies, see [Creating a Signed URL Using a Custom Policy \(p. 166\)](#).

```
function get_custom_policy_stream_name($video_path, $private_key_filename, $key_pair_id,
    $policy) {
    // sign the policy
```



```
$signature = rsa_shal_sign($policy, $private_key_filename);  
// make the signature safe to be included in a url  
$encoded_signature = url_safe_base64_encode($signature);  
  
// combine the above into a stream name  
$stream_name = create_stream_name($video_path, $encoded_policy, $encoded_signature,  
$key_pair_id, null);  
// url-encode the query string characters to work around a flash player bug  
return encode_query_params($stream_name);  
}
```

Full Code Example

The following example code provides a complete demonstration of creating CloudFront signed URLs with PHP. You can download this full example at <https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/samples/demo-php.zip>.

```
<?php  
  
function rsa_shal_sign($policy, $private_key_filename) {  
    $signature = "";  
  
    // load the private key  
    $fp = fopen($private_key_filename, "r");  
    $priv_key = fread($fp, 8192);  
    fclose($fp);  
    $pkeyid = openssl_get_privatekey($priv_key);  
  
    // compute signature  
    openssl_sign($policy, $signature, $pkeyid);  
  
    // free the key from memory  
    openssl_free_key($pkeyid);  
  
    return $signature;  
}  
  
function url_safe_base64_encode($value) {  
    $encoded = base64_encode($value);  
    // replace unsafe characters +, = and / with the safe characters -, _ and ~  
    return str_replace(  
        array('+', '=', '/'),  
        array('-', '_', '~'),  
        $encoded);  
}  
  
function create_stream_name($stream, $policy, $signature, $key_pair_id, $expires) {  
    $result = $stream;  
    // if the stream already contains query parameters, attach the new query parameters to  
    the end  
    // otherwise, add the query parameters  
    $separator = strpos($stream, '?') == FALSE ? '?' : '&';  
    // the presence of an expires time means we're using a canned policy  
    if($expires) {  
        $result .= $path . $separator . "Expires=" . $expires . "&Signature=" .  
$signature . "&Key-Pair-Id=" . $key_pair_id;  
    }  
    // not using a canned policy, include the policy itself in the stream name  
    else {  
        $result .= $path . $separator . "Policy=" . $policy . "&Signature=" . $signature .  
"&Key-Pair-Id=" . $key_pair_id;  
    }  
  
    // new lines would break us, so remove them
```

```

    return str_replace('\n', '', $result);
}

function encode_query_params($stream_name) {
    // Adobe Flash Player has trouble with query parameters being passed into it,
    // so replace the bad characters with their URL-encoded forms
    return str_replace(
        array('?', '=', '&'),
        array('%3F', '%3D', '%26'),
        $stream_name);
}

function get_canned_policy_stream_name($video_path, $private_key_filename, $key_pair_id,
    $expires) {
    // this policy is well known by CloudFront, but you still need to sign it, since it
    // contains your parameters
    $canned_policy = '{"Statement":[{"Resource":"' . $video_path . '","Condition":{
{"DateLessThan":{"AWS:EpochTime":' . $expires . '}}}}]';
    // the policy contains characters that cannot be part of a URL, so we base64 encode it
    $encoded_policy = url_safe_base64_encode($canned_policy);
    // sign the original policy, not the encoded version
    $signature = rsa_shal_sign($canned_policy, $private_key_filename);
    // make the signature safe to be included in a URL
    $encoded_signature = url_safe_base64_encode($signature);

    // combine the above into a stream name
    $stream_name = create_stream_name($video_path, null, $encoded_signature, $key_pair_id,
    $expires);
    // URL-encode the query string characters to support Flash Player
    return encode_query_params($stream_name);
}

function get_custom_policy_stream_name($video_path, $private_key_filename, $key_pair_id,
    $policy) {
    // the policy contains characters that cannot be part of a URL, so we base64 encode it
    $encoded_policy = url_safe_base64_encode($policy);
    // sign the original policy, not the encoded version
    $signature = rsa_shal_sign($policy, $private_key_filename);
    // make the signature safe to be included in a URL
    $encoded_signature = url_safe_base64_encode($signature);

    // combine the above into a stream name
    $stream_name = create_stream_name($video_path, $encoded_policy, $encoded_signature,
    $key_pair_id, null);
    // URL-encode the query string characters to support Flash Player
    return encode_query_params($stream_name);
}

// Path to your private key. Be very careful that this file is not accessible
// from the web!

$private_key_filename = '/home/test/secure/example-priv-key.pem';
$key_pair_id = 'AKIAIOSFODNN7EXAMPLE';

$video_path = 'example.mp4';

$expires = time() + 300; // 5 min from now
$canned_policy_stream_name = get_canned_policy_stream_name($video_path,
    $private_key_filename, $key_pair_id, $expires);

$client_ip = $_SERVER['REMOTE_ADDR'];
$policy =
'{' .
    '"Statement":[' .
        '{' .

```

```

        "Resource": "'. $video_path . "','.
        "Condition":{'.
            "IpAddress":{"AWS:SourceIp":"' . $client_ip . '/32"},'.
            "DateLessThan":{"AWS:EpochTime":"' . $expires . '}'.
        }'.
    }'.
    ']' .
    '}'.
};
$custom_policy_stream_name = get_custom_policy_stream_name($video_path,
    $private_key_filename, $key_pair_id, $policy);
?>

<html>

<head>
    <title>CloudFront</title>
<script type='text/javascript' src='https://example.cloudfront.net/player/swfobject.js'></script>
</head>

<body>
    <h1>Amazon CloudFront</h1>
    <h2>Canned Policy</h2>
    <h3>Expires at <?= gmdate('Y-m-d H:i:s T', $expires) ?></h3>
    <br />

    <div id='canned'>The canned policy video will be here</div>

    <h2>Canned Policy</h2>
    <h3>Expires at <?= gmdate('Y-m-d H:i:s T', $expires) ?> only viewable by IP <?=
    $client_ip ?></h3>
    <div id='custom'>The custom policy video will be here</div>

    <!-- ***** Have to update the player.swf path to a real JWPlayer instance.
    The fake one means that external people cannot watch the video right now -->
    <script type='text/javascript'>
        var so_canned = new SWFObject('https://files.example.com/
        player.swf','mpl','640','360','9');
        so_canned.addParam('allowfullscreen','true');
        so_canned.addParam('allowscriptaccess','always');
        so_canned.addParam('wmode','opaque');
        so_canned.addVariable('file','<?= $canned_policy_stream_name ?>');
        so_canned.addVariable('streamer','rtmp://example.cloudfront.net/cfx/st');
        so_canned.write('canned');

        var so_custom = new SWFObject('https://files.example.com/
        player.swf','mpl','640','360','9');
        so_custom.addParam('allowfullscreen','true');
        so_custom.addParam('allowscriptaccess','always');
        so_custom.addParam('wmode','opaque');
        so_custom.addVariable('file','<?= $custom_policy_stream_name ?>');
        so_custom.addVariable('streamer','rtmp://example.cloudfront.net/cfx/st');
        so_custom.write('custom');
    </script>
</body>

</html>

```

See also:

- [Create a URL Signature Using Perl \(p. 190\)](#)
- [Create a URL Signature Using C# and the .NET Framework \(p. 203\)](#)
- [Create a URL Signature Using Java \(p. 208\)](#)

Create a URL Signature Using C# and the .NET Framework

The C# examples in this section implement an example application that demonstrates how to create the signatures for CloudFront private distributions using canned and custom policy statements. The examples includes utility functions based on the [AWS .NET SDK](#) that can be useful in .NET applications.

You can also create signed URLs and signed cookies by using the AWS SDK for .NET. In the [AWS SDK for .NET API Reference](#), see the following topics:

- **Signed URLs** – Amazon.CloudFront > AmazonCloudFrontUrlSigner
- **Signed cookies** – Amazon.CloudFront > AmazonCloudFrontCookieSigner

Note

Creating a URL signature is just one part of the process of serving private content using a signed URL. For more information about the entire process, see [Using Signed URLs \(p. 157\)](#).

To download the code, go to [Signature Code in C#](#).

To use the RSA keys provided by [AWS Account/Security](#) in the .NET framework, you must convert the AWS-supplied .pem files to the XML format that the .NET framework uses.

After conversion, the RSA private key file is in the following format:

Example RSA Private Key in the XML .NET Framework Format

```
<RSAKeyValue>
  <Modulus>
    wO5IvYCP5UcoCKDo1dcspoMehWBZcyfs9QEzGi6Oe5y+ewGr1oW+vB2GPB
    ANBiVPcUHTFWHwaIBd3oglmF0lGQ1jP/jOfmXHUK2kUUnLnJp+oOBL2NiuFtqcW6h/L5lIpD8Yq+NRHg
    Ty4zDsyr2880MvXv88yEFURCKqEXAMPLE=
  </Modulus>
  <Exponent>AQAB</Exponent>
</P>
  5bmKDaTz
  npENGvqz4Cea8XPH+sxt+2VaAwYnsarVUoSBvT8WlloVuZGG9IZYmH5KteXEufZveYd9UEXAMPLE==
</P>
<Q>
  1v9l/WN1a1N3rOK4VGOCokx7kR2SyTMSbZgF9IWJNOugR/WZw7HTnjipO3c9dy1Ms9pUKwUF4
  6d7049EXAMPLE==
</Q>
<DP>
  RgrSKuLWXMyBH+/l1Dx/I4tXuAJIrlPyo+VmiOc7b5NzHptkSHEPfR9s1
  OK0VqjknclqCJ3Ig86OMEtEXAMPLE==
</DP>
<DQ>
  pjPjvSfw+RoaTu0pgCA/jwW/FGyfN6iim1RFbkT4
  z49DZb2IM885f3vf35eLTaEYRYUHqGZtChNEV0TEXAMPLE==
</DQ>
<InverseQ>
  nkVOJTg5QtGNgWb9i
  cVtzrL/lpFEOHbJXwEJdU99N+7sMK+1066DL/HSBUCD63qD4USpnf0myc24in0EXAMPLE==</InverseQ>
<D>
  Bc7mp7XYHynuPZxChjWNJZiQ+A73gm0ASDv6At7F8Vi9r0xU1Qe/v0AQs3ycN8QlyR4XmbzMLYk
  3yJxFDXo4ZKQtOGzLGteCU2srANiLv26/imXA8FvidZftTAtLviWQZBVPteYIA69ATUYPEq0a5u5wjGy
  UOij9OWyuEXAMPLE=
</D>
</RSAKeyValue>
```

The following C# code creates a signed URL that uses a canned policy by doing the following:

- Creates a policy statement.
- Hashes the policy statement using SHA1, and signs the result using RSA and the private key for your AWS account or for a trusted AWS account that you specify.
- Base64-encodes the hashed and signed policy statement and replaces special characters to make the string safe to use as a URL request parameter.
- Concatenates the values.

For the complete implementation, see the example at [Signature Code in C#](#).

Example Canned Policy Signing Method in C#

```
public static string ToUrlSafeBase64String(byte[] bytes)
{
    return System.Convert.ToBase64String(bytes)
        .Replace('+', '-')
        .Replace('=', '_')
        .Replace('/', '~');
}

public static string CreateCannedPrivateURL(string urlString,
    string durationUnits, string durationNumber, string pathToPolicyStmnt,
    string pathToPrivateKey, string privateKeyId)
{
    // args[] 0-thisMethod, 1-resourceUrl, 2-seconds-minutes-hours-days
    // to expiration, 3-numberOfPreviousUnits, 4-pathToPolicyStmnt,
    // 5-pathToPrivateKey, 6-PrivateKeyId

    TimeSpan timeSpanInterval = GetDuration(durationUnits, durationNumber);

    // Create the policy statement.
    string strPolicy = CreatePolicyStatement(pathToPolicyStmnt,
        urlString,
        DateTime.Now,
        DateTime.Now.Add(timeSpanInterval),
        "0.0.0.0/0");
    if ("Error!" == strPolicy) return "Invalid time frame." +
        "Start time cannot be greater than end time.";

    // Copy the expiration time defined by policy statement.
    string strExpiration = CopyExpirationTimeFromPolicy(strPolicy);

    // Read the policy into a byte buffer.
    byte[] bufferPolicy = Encoding.ASCII.GetBytes(strPolicy);

    // Initialize the SHA1CryptoServiceProvider object and hash the policy data.
    using (SHA1CryptoServiceProvider
        cryptoSHA1 = new SHA1CryptoServiceProvider())
    {
        bufferPolicy = cryptoSHA1.ComputeHash(bufferPolicy);

        // Initialize the RSACryptoServiceProvider object.
        RSACryptoServiceProvider providerRSA = new RSACryptoServiceProvider();
        XmlDocument xmlPrivateKey = new XmlDocument();

        // Load PrivateKey.xml, which you created by converting your
        // .pem file to the XML format that the .NET framework uses.
        // Several tools are available.
        xmlPrivateKey.Load(pathToPrivateKey);

        // Format the RSACryptoServiceProvider providerRSA and
        // create the signature.
```

```
providerRSA.FromXmlString(xmlPrivateKey.InnerXml);
RSAPKCS1SignatureFormatter rsaFormatter =
    new RSAPKCS1SignatureFormatter(providerRSA);
rsaFormatter.SetHashAlgorithm("SHA1");
byte[] signedPolicyHash = rsaFormatter.CreateSignature(bufferPolicy);

// Convert the signed policy to URL-safe base64 encoding and
// replace unsafe characters + = / with the safe characters - _ ~
string strSignedPolicy = ToUrlSafeBase64String(signedPolicyHash);

// Concatenate the URL, the timestamp, the signature,
// and the key pair ID to form the signed URL.
return urlString +
    "?Expires=" +
    strExpiration +
    "&Signature=" +
    strSignedPolicy +
    "&Key-Pair-Id=" +
    privateKeyId;
}
}
```

The following C# code creates a signed URL that uses a custom policy by doing the following:

1. Creates a policy statement.
2. Base64-encodes the policy statement and replaces special characters to make the string safe to use as a URL request parameter.
3. Hashes the policy statement using SHA1, and encrypts the result using RSA and the private key for your AWS account or for a trusted AWS account that you specify.
4. Base64-encodes the hashed policy statement and replacing special characters to make the string safe to use as a URL request parameter.
5. Concatenates the values.

For the complete implementation, see the example at [Signature Code in C#](#).

Example Custom Policy Signing Method in C#

```
public static string ToUrlSafeBase64String(byte[] bytes)
{
    return System.Convert.ToBase64String(bytes)
        .Replace('+', '-')
        .Replace('=', '_')
        .Replace('/', '~');
}

public static string CreateCustomPrivateURL(string urlString,
    string durationUnits, string durationNumber, string startIntervalFromNow,
    string ipaddress, string pathToPolicyStmnt, string pathToPrivateKey,
    string PrivateKeyId)
{
    // args[] 0-thisMethod, 1-resourceUrl, 2-seconds-minutes-hours-days
    // to expiration, 3-numberOfPreviousUnits, 4-starttimeFromNow,
    // 5-ip_address, 6-pathToPolicyStmnt, 7-pathToPrivateKey, 8-privateKeyId

    TimeSpan timeSpanInterval = GetDuration(durationUnits, durationNumber);
    TimeSpan timeSpanToStart = GetDurationByUnits(durationUnits,
        startIntervalFromNow);
    if (null == timeSpanToStart)
        return "Invalid duration units." +
            "Valid options: seconds, minutes, hours, or days";
}
```

```
string strPolicy = CreatePolicyStatement(
    pathToPolicyStmnt, urlString, DateTime.Now.Add(timeSpanToStart),
    DateTime.Now.Add(timeSpanInterval), ipAddress);

// Read the policy into a byte buffer.
byte[] bufferPolicy = Encoding.ASCII.GetBytes(strPolicy);

// Convert the policy statement to URL-safe base64 encoding and
// replace unsafe characters + = / with the safe characters - _ ~

string urlSafePolicy = ToUrlSafeBase64String(bufferPolicy);

// Initialize the SHA1CryptoServiceProvider object and hash the policy data.
byte[] bufferPolicyHash;
using (SHA1CryptoServiceProvider cryptoSHA1 =
    new SHA1CryptoServiceProvider())
{
    bufferPolicyHash = cryptoSHA1.ComputeHash(bufferPolicy);

    // Initialize the RSACryptoServiceProvider object.
    RSACryptoServiceProvider providerRSA = new RSACryptoServiceProvider();
    XmlDocument xmlPrivateKey = new XmlDocument();

    // Load PrivateKey.xml, which you created by converting your
    // .pem file to the XML format that the .NET framework uses.
    // Several tools are available.
    xmlPrivateKey.Load("PrivateKey.xml");

    // Format the RSACryptoServiceProvider providerRSA
    // and create the signature.
    providerRSA.FromXmlString(xmlPrivateKey.InnerXml);
    RSAPKCS1SignatureFormatter RSAFormatter =
        new RSAPKCS1SignatureFormatter(providerRSA);
    RSAFormatter.SetHashAlgorithm("SHA1");
    byte[] signedHash = RSAFormatter.CreateSignature(bufferPolicyHash);

    // Convert the signed policy to URL-safe base64 encoding and
    // replace unsafe characters + = / with the safe characters - _ ~
    string strSignedPolicy = ToUrlSafeBase64String(signedHash);

    return urlString +
        "?Policy=" +
        urlSafePolicy +
        "&Signature=" +
        strSignedPolicy +
        "&Key-Pair-Id=" +
        PrivateKeyId;
}
}
```

Example Utility Methods for Signature Generation

The following methods get the policy statement from a file and parse time intervals for signature generation.

```
public static string CreatePolicyStatement(string policyStmnt,
    string resourceUrl,
    DateTime startTime,
    DateTime endTime,
    string ipAddress)
{
    // Create the policy statement.
```

```

FileStream streamPolicy = new FileStream(policyStmnt, FileMode.Open, FileAccess.Read);
using (StreamReader reader = new StreamReader(streamPolicy))
{
    string strPolicy = reader.ReadToEnd();

    TimeSpan startTimeSpanFromNow = (startTime - DateTime.Now);
    TimeSpan endTimeSpanFromNow = (endTime - DateTime.Now);
    TimeSpan intervalStart =
        (DateTime.UtcNow.Add(startTimeSpanFromNow)) -
        new DateTime(1970, 1, 1, 0, 0, 0, DateTimeKind.Utc);
    TimeSpan intervalEnd =
        (DateTime.UtcNow.Add(endTimeSpanFromNow)) -
        new DateTime(1970, 1, 1, 0, 0, 0, DateTimeKind.Utc);

    int startTimestamp = (int)intervalStart.TotalSeconds; // START_TIME
    int endTimestamp = (int)intervalEnd.TotalSeconds; // END_TIME

    if (startTimestamp > endTimestamp)
        return "Error!";

    // Replace variables in the policy statement.
    strPolicy = strPolicy.Replace("RESOURCE", resourceUrl);
    strPolicy = strPolicy.Replace("START_TIME", startTimestamp.ToString());
    strPolicy = strPolicy.Replace("END_TIME", endTimestamp.ToString());
    strPolicy = strPolicy.Replace("IP_ADDRESS", ipAddress);
    strPolicy = strPolicy.Replace("EXPIRES", endTimestamp.ToString());
    return strPolicy;
}
}

public static TimeSpan GetDuration(string units, string numUnits)
{
    TimeSpan timeSpanInterval = new TimeSpan();
    switch (units)
    {
        case "seconds":
            timeSpanInterval = new TimeSpan(0, 0, 0, int.Parse(numUnits));
            break;
        case "minutes":
            timeSpanInterval = new TimeSpan(0, 0, int.Parse(numUnits), 0);
            break;
        case "hours":
            timeSpanInterval = new TimeSpan(0, int.Parse(numUnits), 0, 0);
            break;
        case "days":
            timeSpanInterval = new TimeSpan(int.Parse(numUnits), 0, 0, 0);
            break;
        default:
            Console.WriteLine("Invalid time units;" +
                "use seconds, minutes, hours, or days");
            break;
    }
    return timeSpanInterval;
}

private static TimeSpan GetDurationByUnits(string durationUnits,
    string startIntervalFromNow)
{
    switch (durationUnits)
    {
        case "seconds":
            return new TimeSpan(0, 0, int.Parse(startIntervalFromNow));
        case "minutes":
            return new TimeSpan(0, int.Parse(startIntervalFromNow), 0);
        case "hours":
            return new TimeSpan(int.Parse(startIntervalFromNow), 0, 0);
    }
}

```



```
        case "days":
            return new TimeSpan(int.Parse(startIntervalFromNow), 0, 0, 0);
        default:
            return new TimeSpan(0, 0, 0, 0);
    }
}

public static string CopyExpirationTimeFromPolicy(string policyStatement)
{
    int startExpiration = policyStatement.IndexOf("EpochTime");
    string strExpirationRough = policyStatement.Substring(startExpiration +
        "EpochTime".Length);
    char[] digits = { '0', '1', '2', '3', '4', '5', '6', '7', '8', '9' };

    List<char> listDigits = new List<char>(digits);
    StringBuilder buildExpiration = new StringBuilder(20);

    foreach (char c in strExpirationRough)
    {
        if (listDigits.Contains(c))
            buildExpiration.Append(c);
    }
    return buildExpiration.ToString();
}
```

See also

- [Create a URL Signature Using Perl \(p. 190\)](#)
- [Create a URL Signature Using PHP \(p. 198\)](#)
- [Create a URL Signature Using Java \(p. 208\)](#)

Create a URL Signature Using Java

In addition to the following code example, you can use [the `CloudFrontUrlSigner` utility class in the AWS SDK for Java \(version 1\)](#) to create [CloudFront signed URLs \(p. 157\)](#).

Note

Creating a signed URL is just one part of the process of [serving private content with CloudFront \(p. 145\)](#). For more information about the entire process, see [Using Signed URLs \(p. 157\)](#).

The following example shows how to create a CloudFront signed URL. You must convert the private key from PEM to DER format for Java implementations to use it.

Example Java Policy and Signature Encryption Methods

```
// Signed URLs for a private distribution
// Note that Java only supports SSL certificates in DER format,
// so you will need to convert your PEM-formatted file to DER format.
// To do this, you can use openssl:
// openssl pkcs8 -topk8 -nocrypt -in origin.pem -inform PEM -out new.der
// -outform DER
// So the encoder works correctly, you should also add the bouncy castle jar
// to your project and then add the provider.

Security.addProvider(new org.bouncycastle.jce.provider.BouncyCastleProvider());

String distributionDomain = "a1b2c3d4e5f6g7.cloudfront.net";
String privateKeyFilePath = "/path/to/rsa-private-key.der";
String s3ObjectKey = "s3/object/key.txt";
```

```
String policyResourcePath = "https://" + distributionDomain + "/" + s3ObjectKey;

// Convert your DER file into a byte array.

byte[] derPrivateKey = ServiceUtils.readInputStreamToBytes(new
    FileInputStream(privateKeyFilePath));

// Generate a "canned" signed URL to allow access to a
// specific distribution and file

String signedUrlCanned = CloudFrontService.signUrlCanned(
    "https://" + distributionDomain + "/" + s3ObjectKey, // Resource URL or Path
    keyPairId, // Certificate identifier,
    // an active trusted signer for the distribution
    derPrivateKey, // DER Private key data
    ServiceUtils.parseIso8601Date("2011-11-14T22:20:00.000Z") // DateLessThan
);
System.out.println(signedUrlCanned);

// Build a policy document to define custom restrictions for a signed URL.

String policy = CloudFrontService.buildPolicyForSignedUrl(
    // Resource path (optional, can include '*' and '?' wildcards)
    policyResourcePath,
    // DateLessThan
    ServiceUtils.parseIso8601Date("2011-11-14T22:20:00.000Z"),
    // CIDR IP address restriction (optional, 0.0.0.0/0 means everyone)
    "0.0.0.0/0",
    // DateGreaterThan (optional)
    ServiceUtils.parseIso8601Date("2011-10-16T06:31:56.000Z")
);

// Generate a signed URL using a custom policy document.

String signedUrl = CloudFrontService.signUrl(
    // Resource URL or Path
    "https://" + distributionDomain + "/" + s3ObjectKey,
    // Certificate identifier, an active trusted signer for the distribution
    keyPairId,
    // DER Private key data
    derPrivateKey,
    // Access control policy
    policy
);
System.out.println(signedUrl);
```

See also:

- [Create a URL Signature Using Perl \(p. 190\)](#)
- [Create a URL Signature Using PHP \(p. 198\)](#)
- [Create a URL Signature Using C# and the .NET Framework \(p. 203\)](#)

Restricting Access to Amazon S3 Content by Using an Origin Access Identity

To restrict access to content that you serve from Amazon S3 buckets, follow these steps:

1. Create a special CloudFront user called an origin access identity (OAI) and associate it with your distribution.

2. Configure your S3 bucket permissions so that CloudFront can use the OAI to access the files in your bucket and serve them to your users. Make sure that users can't use a direct URL to the S3 bucket to access a file there.

After you take these steps, users can only access your files through CloudFront, not directly from the S3 bucket.

In general, if you're using an Amazon S3 bucket as the origin for a CloudFront distribution, you can either allow everyone to have access to the files there, or you can restrict access. If you restrict access by using, for example, CloudFront signed URLs or signed cookies, you also won't want people to be able to view files by simply using the direct Amazon S3 URL for the file. Instead, you want them to only access the files by using the CloudFront URL, so your protections work. For more information about using signed URLs and signed cookies, see [Serving Private Content with Signed URLs and Signed Cookies \(p. 145\)](#).

This topic explains in detail how to set up the OAI and grant permissions to maintain secure access to your S3 files.

Important

If you use an Amazon S3 bucket configured as a website endpoint, you must set it up with CloudFront as a custom origin. You can't use the origin access identity feature described in this topic. However, you *can* restrict access to content on a custom origin by setting up custom headers and configuring your origin to require them. For more information, see [Restricting Access to Files on Custom Origins \(p. 148\)](#).

Topics

- [Overview of OAI Setup \(p. 210\)](#)
- [Creating a CloudFront OAI and Adding it to Your Distribution \(p. 211\)](#)
- [Granting the OAI Permission to Read Files in Your Amazon S3 Bucket \(p. 212\)](#)
- [Using an OAI in Amazon S3 Regions that Support Only Signature Version 4 Authentication \(p. 215\)](#)

Overview of OAI Setup

When you first set up an Amazon S3 bucket as the origin for a CloudFront distribution, you grant everyone permission to read the files in your bucket. This allows anyone to access your files either through CloudFront or using the Amazon S3 URL. CloudFront doesn't expose Amazon S3 URLs, but your users might have those URLs if your application serves any files directly from Amazon S3 or if anyone gives out direct links to specific files in Amazon S3.

If you use CloudFront signed URLs or signed cookies to restrict access to files in your Amazon S3 bucket, you probably also want to prevent users from accessing your Amazon S3 files by using Amazon S3 URLs. If users access your files directly in Amazon S3, they bypass the controls provided by CloudFront signed URLs or signed cookies. This includes control over the date and time that a user can no longer access your content, and control over which IP addresses can be used to access content. In addition, if users access files both through CloudFront and directly by using Amazon S3 URLs, CloudFront access logs are less useful because they're incomplete.

To ensure that your users access your files using only CloudFront URLs, regardless of whether the URLs are signed, do the following:

1. Create an origin access identity, which is a special CloudFront user, and associate the origin access identity with your distribution. You associate the origin access identity with origins, so that you can secure all or just some of your Amazon S3 content. You can also create an origin access identity and add it to your distribution when you create the distribution. For more information, see [Creating a CloudFront OAI and Adding it to Your Distribution \(p. 211\)](#).
2. Change the permissions either on your Amazon S3 bucket or on the files in your bucket so that only the origin access identity has read permission (or read and download permission). When your users

access your Amazon S3 files through CloudFront, the CloudFront origin access identity gets the files on behalf of your users. If your users request files directly by using Amazon S3 URLs, they're denied access. The origin access identity has permission to access files in your Amazon S3 bucket, but users don't. For more information, see [Granting the OAI Permission to Read Files in Your Amazon S3 Bucket](#) (p. 212).

Creating a CloudFront OAI and Adding it to Your Distribution

An AWS account can have [up to 100 CloudFront origin access identities \(OAIs\)](#) (p. 497). However, you can add an OAI to as many distributions as you want, so one OAI is usually sufficient.

If you didn't create an OAI and add it to your distribution when you created the distribution, you can create and add one now by using either the CloudFront console or the CloudFront API:

- **To use the CloudFront console** – You can create an OAI and add it to your distribution at the same time. For step-by-step instructions, see [Creating an OAI and Adding it to Your Distribution](#) (p. 211).
- **To use the CloudFront API** – You create an OAI, and then you add it to your distribution. For step-by-step instructions, see the following:
 - [Creating an OAI Using the CloudFront API](#) (p. 212)
 - [Adding an OAI to Your Distribution Using the CloudFront API](#) (p. 212)

Note

To create OAIs, you must use the CloudFront console, or CloudFront API version 2009-09-09 or later.

Creating an OAI and Adding it to Your Distribution

If you didn't create an OAI when you created your distribution, do the following.

To create a CloudFront OAI using the CloudFront console

1. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
2. Choose the ID of a distribution that has an S3 origin.
3. Choose the **Origins and Origin Groups** tab.
4. Choose the check box next to an origin, and then choose **Edit**.
5. For **Restrict Bucket Access**, choose **Yes**.
6. If you already have an OAI that you want to use, choose **Use an Existing Identity**. Then choose the OAI in the **Your Identities** list.

Note

If you already have an OAI, we recommend that you reuse it to simplify maintenance.

If you want to create an OAI, choose **Create a New Identity**. You can replace the bucket name in the **Comment** field with a custom description.

7. If you want CloudFront to automatically give the OAI permission to read the files in the Amazon S3 bucket specified in **Origin Domain Name**, choose **Yes, Update Bucket Policy**.

Important

If you choose **Yes, Update Bucket Policy**, CloudFront updates bucket permissions to grant the specified OAI permission to read files in your bucket. However, CloudFront does not

remove existing permissions. If users currently have permission to access the files in your bucket using Amazon S3 URLs, they will still have that permission after CloudFront updates your bucket permissions. To view or remove existing bucket permissions, use a method provided by Amazon S3.

If you want to manually update permissions on your Amazon S3 bucket, choose **No, I Will Update Permissions**. For more information, see [Granting the OAI Permission to Read Files in Your Amazon S3 Bucket \(p. 212\)](#).

8. Choose **Yes, Edit**.
9. If you have more than one origin, repeat the steps to add an OAI for each one.

Creating an OAI Using the CloudFront API

If you already have an origin access identity and you want to reuse it instead of creating another one, skip to [Adding an OAI to Your Distribution Using the CloudFront API \(p. 212\)](#).

To create a CloudFront OAI using the CloudFront API, use the `POST Origin Access Identity` API action. The response includes an `Id` and an `S3CanonicalUserId` for the new OAI. Make note of these values because you use them later in the process:

- **Id element** – You use the value of the `Id` element to associate the OAI with your distribution.
- **S3CanonicalUserId element** – You use the value of the `S3CanonicalUserId` element when you use Amazon S3 object ACLs to give the OAI access to your Amazon S3 objects.

For more information, see [CreateCloudFrontOriginAccessIdentity](#) in the *Amazon CloudFront API Reference*.

Adding an OAI to Your Distribution Using the CloudFront API

You can use the CloudFront API to add a CloudFront OAI to an existing distribution or to create a new distribution that includes an OAI. In either case, include an `OriginAccessIdentity` element. This element contains the value of the `Id` element that the `POST Origin Access Identity` API action returned when you created the OAI. You can add the `OriginAccessIdentity` element to one or more origins.

See the following topics in the *Amazon CloudFront API Reference*:

- **Create a new web distribution** – [CreateDistribution](#)
- **Update an existing web distribution** – [UpdateDistribution](#)

Granting the OAI Permission to Read Files in Your Amazon S3 Bucket

When you create or update a distribution, you can add an origin access identity (OAI) and automatically update the Amazon S3 bucket policy to give the OAI permission to access your bucket. Alternatively, you can choose to manually create or update the bucket policy, or use object ACLs that control access to individual files in the bucket.

Whichever method you use, you should still review the permissions to make sure that:

- Your CloudFront OAI can access files in the bucket on behalf of viewers who are requesting them through CloudFront.

- Viewers can't use Amazon S3 URLs to access your files outside of CloudFront.

Important

If you configure CloudFront to accept and forward all of the HTTP methods that CloudFront supports, make sure you give your CloudFront OAI the desired permissions. For example, if you configure CloudFront to accept and forward requests that use the `DELETE` method, configure your bucket policy or object ACLs to handle `DELETE` requests appropriately so viewers can delete only files that you want them to.

Note the following:

- You might find it easier to use Amazon S3 bucket policies than object ACLs because you can add files to the bucket without updating permissions. However, ACLs give you more fine-grained control because you're granting permissions on each individual file.
- By default, your Amazon S3 bucket and all of the files in it are private. Only the AWS account that created the bucket has permission to read or write the files in it.
- If another AWS account uploads files to your bucket, that account is the owner of those files. Bucket policies only apply to files that the bucket owner owns. This means that if another account uploads files to your bucket, the bucket policy that you created for your OAI not evaluated for those files. In that case, use object ACLs to give permissions to your OAI.
- If you're adding an OAI to an existing distribution, modify the bucket policy or any object ACLs as appropriate to ensure that the files are not publicly available outside of CloudFront.
- Grant additional permissions to one or more secure administrator accounts so you can continue to update the contents of the Amazon S3 bucket.

Important

There might be a brief delay between when you save your changes to Amazon S3 permissions and when the changes take effect. Until the changes take effect, you might get "permission denied" errors when you try to access files in your bucket.

Using Amazon S3 Bucket Policies

You can give a CloudFront OAI access to files in an Amazon S3 bucket by creating or updating the bucket policy in the following ways:

- Using the Amazon S3 bucket's **Permissions** tab in the [Amazon S3 console](#).
- Using `PutBucketPolicy` in the Amazon S3 API.
- Using the [CloudFront console](#). When you add an OAI to your origin settings in the CloudFront console, you can choose **Yes, Update Bucket Policy** to tell CloudFront to update the bucket policy on your behalf.

If you update the bucket policy manually, make sure that you:

- Specify the correct OAI as the `Principal` in the policy.
- Give the OAI the permissions it needs to access objects on behalf of viewers.

For more information, see the following sections.

Specify an OAI as the `Principal`

To specify an OAI as the `Principal` in an Amazon S3 bucket policy, use the OAI's Amazon Resource Name (ARN), which includes the OAI's ID. For example:

```
"Principal": {  
  "AWS": "arn:aws:iam::cloudfront:user/CloudFront Origin Access Identity EH1HDMB1FH2TC"  
}
```

To use the preceding example, replace **EH1HDMB1FH2TC** with the OAI's ID. To find the OAI's ID, see the [Origin Access Identity page](#) in the CloudFront console, or use [ListCloudFrontOriginAccessIdentities](#) in the CloudFront API.

You can also specify an OAI as the `Principal` by using its Amazon S3 canonical ID. For example:

```
"Principal": {  
  "CanonicalUser": "79a59df900b949e55d96a1e698fbacedfd6e09d98eacf8f8d5218e7cd47ef2be"  
}
```

To use the preceding example, replace **79a59df900b949e55d96a1e698fbacedfd6e09d98eacf8f8d5218e7cd47ef2be** with the OAI's canonical ID. You can find the OAI's canonical ID in the same ways that you find its ID.

Note

One advantage of using the OAI's ARN to specify it as the `Principal` is that it's easier to understand who the bucket policy is granting access to. Amazon S3 canonical IDs can refer to different kinds of AWS identities, not just CloudFront OAIs, and it can be difficult to determine which identity a canonical ID refers to. Using the OAI's ARN can make it easier to understand the bucket policy.

Also, when you use the OAI's canonical ID in a bucket policy, AWS replaces the canonical ID with the OAI's ARN. When you write a policy that specifies an OAI's canonical ID and then later view the same policy, you'll see that the canonical ID has been replaced by the corresponding ARN. For this reason, it might make sense to just write the policy using the OAI's ARN.

Give Permissions to an OAI

To give the OAI the permissions to access objects in your Amazon S3 bucket, use keywords in the policy that relate to specific Amazon S3 API operations. For example, the `s3:GetObject` permission allows the OAI to read objects in the bucket. For more information, see the examples in the following section, or see [Specifying Permissions in a Policy](#) in the *Amazon Simple Storage Service Developer Guide*.

Amazon S3 Bucket Policy Examples

The following examples show Amazon S3 bucket policies that grant access to a CloudFront OAI. To use these examples:

- Replace **EH1HDMB1FH2TC** with the OAI's ID. To find the OAI's ID, see the [Origin Access Identity page](#) in the CloudFront console, or use [ListCloudFrontOriginAccessIdentities](#) in the CloudFront API.
- Replace **awsexamplebucket** with the name of your Amazon S3 bucket.

Example Amazon S3 bucket policy that gives the OAI read access

The following example allows the OAI to read objects in the specified bucket (`s3:GetObject`).

```
{  
  "Version": "2012-10-17",  
  "Id": "PolicyForCloudFrontPrivateContent",  
  "Statement": [  
    {  
      "Effect": "Allow",  
      "Principal": {  
        "AWS": "arn:aws:iam::cloudfront:user/CloudFront Origin Access  
Identity EH1HDMB1FH2TC"
```

```
    },  
    "Action": "s3:GetObject",  
    "Resource": "arn:aws:s3:::awsexamplebucket/*"  
  }  
]  
}
```

Example Amazon S3 bucket policy that gives the OAI read and write access

The following example allows the OAI to read and write objects in the specified bucket (`s3:GetObject` and `s3:PutObject`). This allows viewers to upload files to your Amazon S3 bucket through CloudFront.

```
{  
  "Version": "2012-10-17",  
  "Id": "PolicyForCloudFrontPrivateContent",  
  "Statement": [  
    {  
      "Effect": "Allow",  
      "Principal": {  
        "AWS": "arn:aws:iam::cloudfront:user/CloudFront Origin Access  
Identity EH1HDMB1FH2TC"  
      },  
      "Action": [  
        "s3:GetObject",  
        "s3:PutObject"  
      ],  
      "Resource": "arn:aws:s3:::aws-example-bucket/*"  
    }  
  ]  
}
```

Updating Amazon S3 Object ACLs

You can give a CloudFront OAI access to files in an Amazon S3 bucket by creating or updating the file's ACL in the following ways:

- Using the Amazon S3 object's **Permissions** tab in the [Amazon S3 console](#).
- Using `PutObjectAcl` in the Amazon S3 API.

When you grant access to an OAI using an ACL, you must specify the OAI using its Amazon S3 canonical user ID. This is the value of **Amazon S3 Canonical User ID** on the [Origin Access Identity](#) page in the CloudFront console. If you're using the CloudFront API, use the value of the `S3CanonicalUserId` element that was returned when you created the OAI, or call [ListCloudFrontOriginAccessIdentities](#) in the CloudFront API.

For more information about Amazon S3 object ACLs, see [Managing Access with ACLs](#) in the *Amazon Simple Storage Service Developer Guide*.

Using an OAI in Amazon S3 Regions that Support Only Signature Version 4 Authentication

Newer Amazon S3 Regions require that you use Signature Version 4 for authenticated requests. (For the versions of signature supported in each Amazon S3 Region, see [Amazon Simple Storage Service \(S3\)](#) in the topic [Regions and Endpoints](#) in the *Amazon Web Services General Reference*.) However, when you create an origin access identity and add it to a CloudFront distribution, CloudFront typically uses Signature Version 4 for authentication when it requests files in your Amazon S3 bucket. If you're using

an origin access identity and if your bucket is in one of the Regions that requires Signature Version 4 for authentication, note the following:

- `DELETE`, `GET`, `HEAD`, `OPTIONS`, and `PATCH` requests are supported without qualifications.
- If you want to submit `PUT` requests to CloudFront to upload files to your Amazon S3 bucket, you must add an `x-amz-content-sha256` header to the request. The header value must contain an SHA256 hash of the body of the request. For more information, see the documentation about the `x-amz-content-sha256` header on the [Common Request Headers](#) page in the *Amazon Simple Storage Service API Reference*.
- `POST` requests are not supported.

Using AWS WAF to Control Access to Your Content

AWS WAF is a web application firewall that lets you monitor the HTTP and HTTPS requests that are forwarded to CloudFront, and lets you control access to your content. Based on conditions that you specify, such as the values of query strings or the IP addresses that requests originate from, CloudFront responds to requests either with the requested content or with an HTTP status code 403 (`Forbidden`). You can also configure CloudFront to return a custom error page when a request is blocked. For more information about AWS WAF, see the [AWS WAF Developer Guide](#).

After you create an AWS WAF web access control list (web ACL), create or update a web distribution to associate the distribution with the web ACL. You can associate as many CloudFront distributions as you want with the same web ACL or with different web ACLs. For information about creating a web distribution and associating it with a web ACL, see [Creating a Distribution](#) (p. 37).

To associate or disassociate a web ACL and an existing distribution, or change the web ACL that is associated with a distribution, perform the following procedure.

To associate or disassociate an AWS WAF web ACL and an existing CloudFront distribution by using the CloudFront console

1. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
2. Choose the ID for the distribution that you want to update.
3. On the **General** tab, choose **Edit**.
4. On the **Distribution Settings** page, in the **AWS WAF Web ACL** list, choose the web ACL that you want to associate with this distribution.

If you want to disassociate the distribution from all web ACLs, choose **None**. If you want to associate the distribution with a different web ACL, choose the new web ACL.

5. Choose **Yes, Edit**.
6. Repeat steps 2 through 5 for other distributions, if any, for which you want to add, delete, or change associations with AWS WAF web ACLs.
7. After you change settings, the value of the **Status** column for the distributions that you updated changes to **InProgress** while CloudFront propagates the changes to edge locations. When **Status** changes to **Deployed** for a distribution, the distribution is ready to use AWS WAF when it processes requests. (The value of the **State** column for the distribution must also be **Enabled**.) This should take less than 15 minutes after you save the last change to a distribution.

Note

[AWS Firewall Manager](#) is a security management service that makes it easier to centrally configure and manage AWS WAF rules across your accounts and applications. Using Firewall Manager, you can roll out AWS WAF rules to your CloudFront distributions across accounts in AWS Organizations. For more information, see the [AWS Firewall Manager Developer Guide](#).

Restricting the Geographic Distribution of Your Content

You can use *geo restriction*, also known as *geo blocking*, to prevent users in specific geographic locations from accessing content that you're distributing through a CloudFront web distribution. To use geo restriction, you have two options:

- Use the CloudFront geo restriction feature. Use this option to restrict access to all of the files that are associated with a distribution and to restrict access at the country level.
- Use a third-party geolocation service. Use this option to restrict access to a subset of the files that are associated with a distribution or to restrict access at a finer granularity than the country level.

Topics

- [Using CloudFront Geo Restriction \(p. 217\)](#)
- [Using a Third-Party Geolocation Service \(p. 218\)](#)

Using CloudFront Geo Restriction

When a user requests your content, CloudFront typically serves the requested content regardless of where the user is located. If you need to prevent users in specific countries from accessing your content, you can use the CloudFront geo restriction feature to do one of the following:

- Allow your users to access your content only if they're in one of the countries on a whitelist of approved countries.
- Prevent your users from accessing your content if they're in one of the countries on a blacklist of banned countries.

For example, if a request comes from a country where, for copyright reasons, you are not authorized to distribute your content, you can use CloudFront geo restriction to block the request.

Note

CloudFront determines the location of your users by using a third-party GeoIP database. The accuracy of the mapping between IP addresses and countries varies by Region. Based on recent tests, the overall accuracy is 99.8%. If CloudFront can't determine a user's location, CloudFront serves the content that the user has requested.

Here's how geo restriction works:

1. Suppose you have rights to distribute your content only in Liechtenstein. You update your CloudFront web distribution and add a whitelist that contains only Liechtenstein. (Alternatively, you could add a blacklist that contains every country except Liechtenstein.)
2. A user in Monaco requests your content, and DNS routes the request to the CloudFront edge location in Milan, Italy.
3. The edge location in Milan looks up your distribution and determines that the user in Monaco is not allowed to download your content.
4. CloudFront returns an HTTP status code 403 (*Forbidden*) to the user.

You can optionally configure CloudFront to return a custom error message to the user, and you can specify how long you want CloudFront to cache the error response for the requested file. The default value is one minute. For more information, see [Creating a Custom Error Page for Specific HTTP Status Codes \(p. 291\)](#).

Geo restriction applies to an entire web distribution. If you need to apply one restriction to part of your content and a different restriction (or no restriction) to another part of your content, you must either create separate CloudFront web distributions or use a third-party geolocation service.

If you enable CloudFront access logging, you can identify the requests that CloudFront rejected by searching for the log entries for which the value of `sc-status` (the HTTP status code) is 403. However, using only the access logs, you can't distinguish a request that CloudFront rejected based on the location of the user from a request that CloudFront rejected because the user didn't have permission to access the file for another reason. If you have a third-party geolocation service such as Digital Element or MaxMind, you can identify the location of requests based on the IP address in the `c-ip` (client IP) column in the access logs. For more information about CloudFront access logs, see [Configuring and Using Standard Logs \(Access Logs\)](#) (p. 439).

The following procedure explains how to use the CloudFront console to add geo restriction to an existing web distribution. For information about how to use the console to create a web distribution, see [Creating a Distribution](#) (p. 37).

To add geo restriction to your CloudFront web distribution (console)

1. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
2. Select the distribution that you want to update.
3. In the **Distribution Settings** pane, choose the **Restrictions** tab.
4. Choose **Edit**.
5. Enter the applicable values. For more information, see [Restrictions](#) (p. 61).
6. Choose **Yes, Edit**.

Using a Third-Party Geolocation Service

The CloudFront geo restriction feature lets you control distribution of your content at the country level for all files that you're distributing with a given web distribution. If you have geographic restrictions on where your content can be distributed and the restrictions don't follow country boundaries, or if you want to restrict access to only some of the files that you're distributing through CloudFront, you can combine CloudFront with a third-party geolocation service. This can allow you to control access to your content based not only on country but also based on city, zip or postal code, or even latitude and longitude.

When you're using a third-party geolocation service, we recommend that you use CloudFront signed URLs, which let you specify an expiration date and time after which the URL is no longer valid. In addition, we recommend that you use an Amazon S3 bucket as your origin because you can then use a CloudFront origin access identity to prevent users from accessing your content directly from the origin. For more information about signed URLs and origin access identities, see [Serving Private Content with Signed URLs and Signed Cookies](#) (p. 145).

The following steps explain how to control access to your files by using a third-party geolocation service.

To use geographic location to restrict access to files in a CloudFront distribution

1. Get an account with a geolocation service.
2. Upload your content to an Amazon Simple Storage Service (S3) bucket. For more information, see the [Amazon S3 documentation](#).
3. Configure Amazon CloudFront and Amazon S3 to serve private content. For more information, see [Serving Private Content with Signed URLs and Signed Cookies](#) (p. 145).
4. Write your web application to do the following:

- a. Send the IP address for each user request to the geolocation service.
- b. Evaluate the return value from the geolocation service to determine whether the user is in a location to which you want CloudFront to distribute your content.
- c. If you want to distribute your content to the user's location, generate a signed URL for your CloudFront content. If you don't want to distribute content to that location, return HTTP status code 403 (`Forbidden`) to the user. Alternatively, you can configure CloudFront to return a custom error message. For more information, see [Creating a Custom Error Page for Specific HTTP Status Codes \(p. 291\)](#).

For more information, refer to the documentation for the geolocation service that you're using.

You can use a web server variable to get the IP addresses of the users who are visiting your website. Note the following caveats:

- If your web server is not connected to the internet through a load balancer, you can use a web server variable to get the remote IP address. However, this IP address isn't always the user's IP address. It can also be the IP address of a proxy server, depending on how the user is connected to the internet.
- If your web server is connected to the internet through a load balancer, a web server variable might contain the IP address of the load balancer, not the IP address of the user. In this configuration, we recommend that you use the last IP address in the `X-Forwarded-For` HTTP header. This header typically contains more than one IP address, most of which are for proxies or load balancers. The last IP address in the list is the one most likely to be associated with the user's geographic location.

If your web server is not connected to a load balancer, we recommend that you use web server variables instead of the `X-Forwarded-For` header to avoid IP address spoofing.

Using Field-Level Encryption to Help Protect Sensitive Data

With Amazon CloudFront, you can enforce secure end-to-end connections to origin servers by using HTTPS. Field-level encryption adds an additional layer of security that lets you protect specific data throughout system processing so that only certain applications can see it.

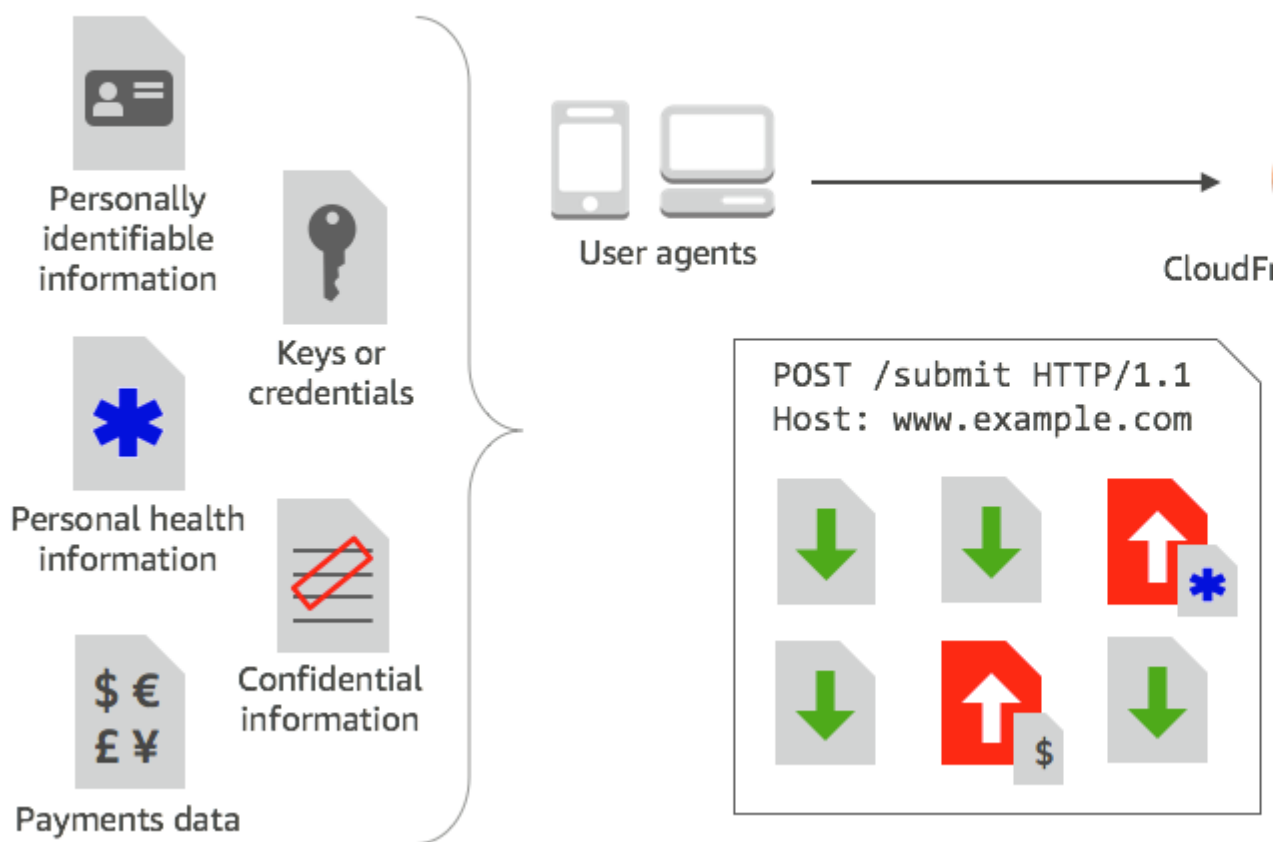
Field-level encryption allows you to enable your users to securely upload sensitive information to your web servers. The sensitive information provided by your users is encrypted at the edge, close to the user, and remains encrypted throughout your entire application stack. This encryption ensures that only applications that need the data—and have the credentials to decrypt it—are able to do so.

To use field-level encryption, when you configure your CloudFront distribution, specify the set of fields in POST requests that you want to be encrypted, and the public key to use to encrypt them. You can encrypt up to 10 data fields in a request. (You can't encrypt all of the data in a request with field-level encryption; you must specify individual fields to encrypt.)

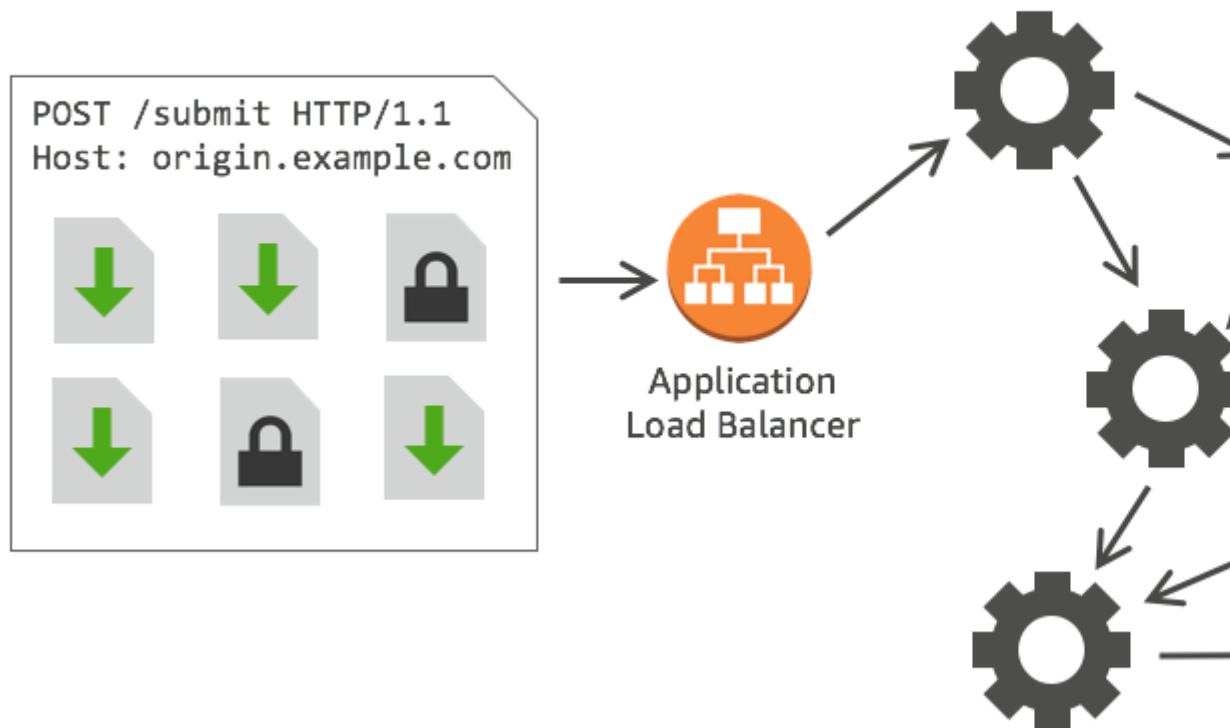
When the HTTPS request with field-level encryption is forwarded to the origin, and the request is routed throughout your origin application or subsystem, the sensitive data is still encrypted, reducing the risk of a data breach or accidental data loss of the sensitive data. Components that need access to the sensitive data for business reasons, such as a payment processing system needing access to a credit number, can use the appropriate private key to decrypt and access the data.

Note

To use field-level encryption, your origin must support chunked encoding.



CloudFront field-level encryption uses asymmetric encryption, also known as public key encryption. You provide a public key to CloudFront, and all sensitive data that you specify is encrypted automatically. The key you provide to CloudFront cannot be used to decrypt the encrypted values; only your private key can do that.



Topics

- [Overview of Field-Level Encryption \(p. 221\)](#)
- [Setting Up Field-Level Encryption \(p. 222\)](#)
- [Decrypting Data Fields at Your Origin \(p. 225\)](#)

Overview of Field-Level Encryption

The following steps provide an overview of setting up field-level encryption. For specific steps, see [Setting Up Field-Level Encryption \(p. 222\)](#).

1. **Get a public key-private key pair.** You must obtain and add the public key before you start setting up field-level encryption in CloudFront.
2. **Create a field-level encryption profile.** Field-level encryption profiles, which you create in CloudFront, define the fields that you want to be encrypted.
3. **Create a field-level encryption configuration.** A configuration specifies the profiles to use, based on the content type of the request or a query argument, for encrypting specific data fields. You can also choose the request-forwarding behavior options that you want for different scenarios. For example, you can set the behavior for when the profile name specified by the query argument in a request URL doesn't exist in CloudFront.

4. **Link to a cache behavior.** Link the configuration to a cache behavior for a distribution, to specify when CloudFront should encrypt data.

Setting Up Field-Level Encryption

Follow these steps to get started using field-level encryption. To learn about quotas (formerly known as limits) on field-level encryption, see [Quotas](#) (p. 496).

- [Step 1: Create an RSA Key Pair](#) (p. 222)
- [Step 2: Add Your Public Key to CloudFront](#) (p. 222)
- [Step 3: Create a Profile for Field-Level Encryption](#) (p. 223)
- [Step 4: Create a Configuration](#) (p. 223)
- [Step 5: Add a Configuration to a Cache Behavior](#) (p. 225)

Step 1: Create an RSA Key Pair

To get started, you must create an RSA key pair that includes a public key and a private key. The public key enables CloudFront to encrypt data, and the private key enables components at your origin to decrypt the fields that have been encrypted. You can use OpenSSL or another tool to create a key pair. The key size must be 2048 bits.

For example, if you're using OpenSSL, you can use the following command to generate a key pair with a length of 2048 bits and save it in the file `private_key.pem`:

```
openssl genrsa -out private_key.pem 2048
```

The resulting file contains both the public and the private key. To extract the public key from that file, run the following command:

```
openssl rsa -pubout -in private_key.pem -out public_key.pem
```

The public key file (`public_key.pem`) contains the encoded key value that you paste in the following step.

Step 2: Add Your Public Key to CloudFront

After you get your RSA key pair, add your public key to CloudFront.

To add your public key to CloudFront (console)

1. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
2. In the navigation pane, choose **Public key**.
3. Choose **Add public key**.
4. For **Key name**, type a unique name for the key. The name can't have spaces and can include only alphanumeric characters, underscores (_), and hyphens (-). The maximum number of characters is 128.
5. For **Key value**, paste the encoded key value for your public key, including the -----BEGIN PUBLIC KEY----- and -----END PUBLIC KEY----- lines.
6. For **Comment**, add an optional comment. For example, you could include the expiration date for the public key.
7. Choose **Add**.

You can add more keys to use with CloudFront by repeating the steps in the procedure.

Step 3: Create a Profile for Field-Level Encryption

After you add at least one public key to CloudFront, create a profile that tells CloudFront which fields to encrypt.

To create a profile for field-level encryption (console)

1. In the navigation pane, choose **Field-level encryption**.
2. Choose **Create profile**.
3. Fill in the following fields:

Profile name

Type a unique name for the profile. The name can't have spaces and can include only alphanumeric characters, underscores (_), and hyphens (-). The maximum number of characters is 128.

Public key name

In the drop-down list, choose the name of a public key that you added to CloudFront in step 2. CloudFront uses the key to encrypt the fields that you specify in this profile.

Provider name

Type a phrase to help identify the key, such as the provider where you got the key pair. This information, along with the private key, is needed when applications decrypt data fields. The provider name can't have spaces and can include only alphanumeric characters, colons (:), underscores (_), and hyphens (-). The maximum number of characters is 128.

Field name pattern to match

Type the names of the data fields, or patterns that identify data field names in the request, that you want CloudFront to encrypt. Choose the + option to add all the fields that you want to encrypt with this key.

For the field name pattern, you can type the entire name of the data field, like DateOfBirth, or just the first part of the name with a wildcard character (*), like CreditCard*. The field name pattern must include only alphanumeric characters, square brackets ([and]), periods (.), underscores (_), and hyphens (-), in addition to the optional wildcard character (*).

Make sure that you don't use overlapping characters for different field name patterns. For example, if you have a field name pattern of ABC*, you can't add another field name pattern that is AB*. In addition, field names are case-sensitive and the maximum number of characters that you can use is 128.

Comment

(Optional) Type a comment about this profile. The maximum number of characters that you can use is 128.

4. After you fill in the fields, choose **Create profile**.
5. If you want to add more profiles, choose **Add profile**.

Step 4: Create a Configuration

After you create one or more field-level encryption profiles, create a configuration that specifies the content type of the request that includes the data to be encrypted, the profile to use for encryption, and other options that specify how you want CloudFront to handle encryption.

For example, when CloudFront can't encrypt the data, you can specify whether CloudFront should block or forward a request to your origin in the following scenarios:

- **When a request's content type isn't in a configuration** – If you haven't added a content type to a configuration, you can specify whether CloudFront should forward the request with that content type to the origin without encrypting data fields, or block the request and return an error.

Note

If you add a content type to a configuration but haven't specified a profile to use with that type, CloudFront always forwards requests with that content type to the origin.

- **When the profile name provided in a query argument is unknown** – When you specify the `file-profile` query argument with a profile name that doesn't exist for your distribution, you can specify whether CloudFront should send the request to the origin without encrypting data fields, or block the request and return an error.

In a configuration, you can also specify whether providing a profile as a query argument in a URL overrides a profile that you've mapped to the content type for that query. By default, CloudFront uses the profile that you've mapped to a content type, if you specify one. This lets you have a profile that's used by default but decide for certain requests that you want to enforce a different profile.

So, for example, you might specify (in your configuration) **SampleProfile** as the query argument profile to use. Then you could use the URL `https://d1234.cloudfront.net?file-profile=SampleProfile` instead of `https://d1234.cloudfront.net`, to have CloudFront use **SampleProfile** for this request, instead of the profile you'd set up for the content type of the request.

You can create up to 10 configurations for a single account, and then associate one of the configurations to the cache behavior of any distribution for the account.

To create a configuration for field-level encryption (console)

1. On the **Field-level encryption** page, choose **Create configuration**.

Note: If you haven't created at least one profile, you won't see the option to create a configuration.

2. Fill in the following fields to specify the profile to use. (Some fields can't be changed.)

Content type (can't be changed)

The content type is set to `application/x-www-form-urlencoded` and can't be changed.

Default profile ID (optional)

In the drop-down list, choose the profile that you want to map to the content type in the **Content type** field.

Content format (can't be changed)

The content format is set to `URLencoded` and can't be changed.

3. If you want to change the CloudFront default behavior for the following options, select the appropriate check box.

Forward request to origin when request's content type is not configured

Select the check box if you want to allow the request to go to your origin *if you have not specified a profile to use for the content type of the request*.

Override the profile for a content type with a provided query argument

Select the check box if you want to allow a profile provided in a query argument to *override the profile that you've specified for a content type*.

4. If you select the check box to allow a query argument to override the default profile, you must complete the following additional fields for the configuration. You can create up to five of these query argument mappings to use with queries.

Query argument

Type the value that you want to include in URLs for the `file-profile` query argument. This value tells CloudFront to use the profile ID (that you specify in the next field) associated with this query argument for field-level encryption for this query.

The maximum number of characters that you can use is 128. The value can't include spaces, and must use only alphanumeric or the following characters: dash (-), period (.), underscore (_), asterisk (*), plus-sign (+), percent (%).

Profile ID

In the drop-down list, choose the profile that you want to associate with the value that you typed for **Query argument**.

Forward request to origin when the profile specified in a query argument does not exist

Select the check box if you want to allow the request to go to your origin *if the profile specified in a query argument isn't defined in CloudFront*.

Step 5: Add a Configuration to a Cache Behavior

To use field-level encryption, link a configuration to a cache behavior for a distribution by adding the configuration ID as a value for your distribution.

Important

To link a field-level encryption configuration to a cache behavior, the distribution must be configured to always use HTTPS, and to accept HTTP `POST` and `PUT` requests from viewers. That is, the following must be true:

- The cache behavior's **Viewer Protocol Policy** must be set to **Redirect HTTP to HTTPS** or **HTTPS Only**. (In AWS CloudFormation or the CloudFront API, `ViewerProtocolPolicy` must be set to `redirect-to-https` or `https-only`.)
- The cache behavior's **Allowed HTTP Methods** must be set to **GET, HEAD, OPTIONS, PUT, POST, PATCH, DELETE**. (In AWS CloudFormation or the CloudFront API, `AllowedMethods` must be set to `GET, HEAD, OPTIONS, PUT, POST, PATCH, DELETE`. These can be specified in any order.)
- The origin setting's **Origin Protocol Policy** must be set to **Match Viewer** or **HTTPS Only**. (In AWS CloudFormation or the CloudFront API, `OriginProtocolPolicy` must be set to `match-viewer` or `https-only`.)

For more information, see [Values That You Specify When You Create or Update a Distribution \(p. 38\)](#).

Decrypting Data Fields at Your Origin

CloudFront encrypts data fields by using the [AWS Encryption SDK](#). The data remains encrypted throughout your application stack and can be accessed only by applications that have the credentials to decrypt it.

After encryption, the ciphertext is base64 encoded. When your applications decrypt the text at the origin, they must first decode the ciphertext, and then use the AWS Encryption SDK to decrypt the data.

The following code example illustrates how applications can decrypt data at your origin. Note the following:

- To simplify the example, this sample loads public and private keys (in DER format) from files in the working directory. In practice, you would store the private key in a secure offline location, such as an offline hardware security module, and distribute the public key to your development team.
- CloudFront uses specific information while encrypting the data, and the same set of parameters should be used at the origin to decrypt it. Parameters CloudFront uses while initializing the MasterKey include the following:
 - PROVIDER_NAME: You specified this value when you created a field-level encryption profile. Use the same value here.
 - KEY_NAME: You created a name for your public key when you uploaded it to CloudFront, and then specified the key name in the profile. Use the same value here.
 - ALGORITHM: CloudFront uses RSA/ECB/OAEPWithSHA-256AndMGF1Padding as the algorithm for encrypting, so you must use the same algorithm to decrypt the data.
- If you run the following sample program with ciphertext as input, the decrypted data is output to your console. For more information, see the [Java Example Code](#) in the AWS Encryption SDK.

Sample Code

```
import java.nio.file.Files;
import java.nio.file.Paths;
import java.security.KeyFactory;
import java.security.PrivateKey;
import java.security.PublicKey;
import java.security.spec.PKCS8EncodedKeySpec;
import java.security.spec.X509EncodedKeySpec;

import org.apache.commons.codec.binary.Base64;

import com.amazonaws.encryptionsdk.AwsCrypto;
import com.amazonaws.encryptionsdk.CryptoResult;
import com.amazonaws.encryptionsdk.jce.JceMasterKey;

/**
 * Sample example of decrypting data that has been encrypted by CloudFront Field-Level
 * Encryption.
 */
public class DecryptExample {

    private static final String PRIVATE_KEY_FILENAME = "private_key.der";
    private static final String PUBLIC_KEY_FILENAME = "public_key.der";
    private static PublicKey publicKey;
    private static PrivateKey privateKey;

    // CloudFront uses the following values to encrypt data, and your origin must use same
    // values to decrypt it.
    // In your own code, for PROVIDER_NAME, use the provider name that you specified when
    // you created your Field Level
    // Encryption Profile. This sample uses 'DEMO' for the value.
    private static final String PROVIDER_NAME = "DEMO";
    // In your own code, use the Key name that you specified when you added your public key
    // to CloudFront. This sample
    // uses 'DEMOKEY' for the Key name.
    private static final String KEY_NAME = "DEMOKEY";
    // Cloudfront uses this algorithm when encrypting data.
    private static final String ALGORITHM = "RSA/ECB/OAEPWithSHA-256AndMGF1Padding";

    public static void main(final String[] args) throws Exception {

        final String dataToDecrypt = args[0];
```

```
// This sample uses files to get public and private keys.
// In practice, you should distribute the public key and save the private key in
secure storage.
populateKeyPair();

System.out.println(decrypt(debase64(dataToDecrypt)));
}

private static String decrypt(final byte[] bytesToDecrypt) throws Exception {
    // You can decrypt the stream only by using the private key.

    // 1. Instantiate the SDK
    final AwsCrypto crypto = new AwsCrypto();

    // 2. Instantiate a JCE master key
    final JceMasterKey masterKey = JceMasterKey.getInstance(
        publicKey,
        privateKey,
        PROVIDER_NAME,
        KEY_NAME,
        ALGORITHM);

    // 3. Decrypt the data
    final CryptoResult <byte[], ? > result = crypto.decryptData(masterKey,
bytesToDecrypt);
    return new String(result.getResult());
}

// Function to decode base64 cipher text.
private static byte[] debase64(final String value) {
    return Base64.decodeBase64(value.getBytes());
}

private static void populateKeyPair() throws Exception {
    final byte[] PublicKeyBytes = Files.readAllBytes(Paths.get(PUBLIC_KEY_FILENAME));
    final byte[] privateKeyBytes = Files.readAllBytes(Paths.get(PRIVATE_KEY_FILENAME));
    publicKey = KeyFactory.getInstance("RSA").generatePublic(new
X509EncodedKeySpec(PublicKeyBytes));
    privateKey = KeyFactory.getInstance("RSA").generatePrivate(new
PKCS8EncodedKeySpec(privateKeyBytes));
}
}
```

Optimizing Content Caching and Availability

This section describes how to set up and manage the caching of objects to improve performance and meet your business requirements.

To learn about adding and removing the content that you want CloudFront to serve, see [Adding, Removing, or Replacing Content That CloudFront Distributes](#) (p. 104).

Topics

- [How Caching Works with CloudFront Edge Caches](#) (p. 228)
- [Increasing the Proportion of Requests that Are Served from CloudFront Edge Caches \(Cache Hit Ratio\)](#) (p. 228)
- [Optimizing High Availability with CloudFront Origin Failover](#) (p. 231)
- [Managing How Long Content Stays in an Edge Cache \(Expiration\)](#) (p. 236)
- [Caching Content Based on Query String Parameters](#) (p. 241)
- [Caching Content Based on Cookies](#) (p. 244)
- [Caching Content Based on Request Headers](#) (p. 246)

How Caching Works with CloudFront Edge Caches

One of the purposes of using CloudFront is to reduce the number of requests that your origin server must respond to directly. With CloudFront caching, more objects are served from CloudFront edge locations, which are closer to your users. This reduces the load on your origin server and reduces latency.

The more requests that CloudFront can serve from edge caches, the fewer viewer requests that CloudFront must forward to your origin to get the latest version or a unique version of an object.

The proportion of requests that are served from caches to all requests is called the *cache hit ratio*. You can view the percentage of viewer requests that are hits, misses, and errors in the CloudFront console. For more information, see [CloudFront Cache Statistics Reports](#) (p. 407).

A number of factors affect the cache hit ratio. You can adjust your CloudFront distribution configuration to improve the cache hit ratio by following the guidance in [Increasing the Proportion of Requests that Are Served from CloudFront Edge Caches \(Cache Hit Ratio\)](#) (p. 228).

Increasing the Proportion of Requests that Are Served from CloudFront Edge Caches (Cache Hit Ratio)

You can improve performance by increasing the proportion of your viewer requests that are served from CloudFront edge caches instead of going to your origin servers for content. This is known as improving the cache hit ratio for your distribution.

The following sections explain how to improve your cache hit ratio.

Topics

- [Specifying How Long CloudFront Caches Your Objects](#) (p. 229)
- [Caching Based on Query String Parameters](#) (p. 229)
- [Caching Based on Cookie Values](#) (p. 229)
- [Caching Based on Request Headers](#) (p. 230)
- [Remove Accept-Encoding Header When Compression is Not Needed](#) (p. 231)
- [Serving Media Content by Using HTTP](#) (p. 231)

Specifying How Long CloudFront Caches Your Objects

To increase your cache hit ratio, you can configure your origin to add a `Cache-Control max-age` directive to your objects, and specify the longest practical value for `max-age`. The shorter the cache duration, the more frequently CloudFront forwards requests to your origin to determine if an object has changed and to get the latest version. For more information, see [Managing How Long Content Stays in an Edge Cache \(Expiration\)](#) (p. 236).

Caching Based on Query String Parameters

If you configure CloudFront to cache based on query string parameters, you can improve caching if you do the following:

- Configure CloudFront to forward only the query string parameters for which your origin will return unique objects.
- Use the same case (uppercase or lowercase) for all instances of the same parameter. For example, if one request contains `parameter1=A` and another contains `parameter1=a`, CloudFront forwards separate requests to your origin when a request contains `parameter1=A` and when a request contains `parameter1=a`. CloudFront then separately caches the corresponding objects returned by your origin separately even if the objects are identical. If you use just `A` or `a`, CloudFront forwards fewer requests to your origin.
- List parameters in the same order. As with differences in case, if one request for an object contains the query string `parameter1=a¶meter2=b` and another request for the same object contains `parameter2=b¶meter1=a`, CloudFront forwards both requests to your origin and separately caches the corresponding objects even if they're identical. If you always use the same order for parameters, CloudFront forwards fewer requests to your origin.

For more information, see [Caching Content Based on Query String Parameters](#) (p. 241). If you want to review the query strings that CloudFront forwards to your origin, see the values in the `cs-uri-query` column of your CloudFront log files. For more information, see [Configuring and Using Standard Logs \(Access Logs\)](#) (p. 439).

Caching Based on Cookie Values

If you configure CloudFront to cache based on cookie values, you can improve caching if you do the following:

- Configure CloudFront to forward only specified cookies instead of forwarding all cookies. For the cookies that you configure CloudFront to forward to your origin, CloudFront forwards every combination of cookie name and value. It then separately caches the objects that your origin returns, even if they're all identical.

For example, suppose that viewers include two cookies in every request, that each cookie has three possible values, and that all combinations of cookie values are possible. CloudFront forwards up to six

different requests to your origin for each object. If your origin returns different versions of an object based on only one of the cookies, then CloudFront is forwarding more requests to your origin than necessary and is needlessly caching multiple identical versions of the object.

- Create separate cache behaviors for static and dynamic content, and configure CloudFront to forward cookies to your origin only for dynamic content.

For example, suppose you have just one cache behavior for your distribution and that you're using the distribution both for dynamic content, such as `.js` files, and for `.css` files that rarely change. CloudFront caches separate versions of your `.css` files based on cookie values, so each CloudFront edge location forwards a request to your origin for every new cookie value or combination of cookie values.

If you create a cache behavior for which the path pattern is `*.css` and for which CloudFront doesn't cache based on cookie values, then CloudFront forwards requests for `.css` files to your origin for only the first request that an edge location receives for a given `.css` file and for the first request after a `.css` file expires.

- If possible, create separate cache behaviors for dynamic content when cookie values are unique for each user (such as a user ID), and dynamic content that varies based on a smaller number of unique values.

For more information, see [Caching Content Based on Cookies \(p. 244\)](#). If you want to review the cookies that CloudFront forwards to your origin, see the values in the `cs(Cookie)` column of your CloudFront log files. For more information, see [Configuring and Using Standard Logs \(Access Logs\) \(p. 439\)](#).

Caching Based on Request Headers

If you configure CloudFront to cache based on request headers, you can improve caching if you do the following:

- Configure CloudFront to forward and cache based on only specified headers instead of forwarding and caching based on all headers. For the headers that you specify, CloudFront forwards every combination of header name and value. It then separately caches the objects that your origin returns even if they're all identical.

Note

CloudFront always forwards to your origin the headers specified in the following topics:

- [How CloudFront Processes and Forwards Requests to Your Amazon S3 Origin Server > HTTP Request Headers That CloudFront Removes or Updates \(p. 266\)](#)
- [How CloudFront Processes and Forwards Requests to Your Custom Origin Server > HTTP Request Headers and CloudFront Behavior \(Custom and S3 Origins\) \(p. 273\)](#)

When you configure CloudFront to cache based on request headers, you don't change the headers that CloudFront forwards, only whether CloudFront caches objects based on the header values.

- Try to avoid caching based on request headers that have large numbers of unique values.

For example, if you want to serve different sizes of an image based on the user's device, then don't configure CloudFront to cache based on the `User-Agent` header, which has an enormous number of possible values. Instead, configure CloudFront to cache based on the CloudFront device-type headers `CloudFront-Is-Desktop-Viewer`, `CloudFront-Is-Mobile-Viewer`, `CloudFront-Is-SmartTV-Viewer`, and `CloudFront-Is-Tablet-Viewer`. In addition, if you're returning the same version of the image for tablets and desktops, then forward only the `CloudFront-Is-Tablet-Viewer` header, not the `CloudFront-Is-Desktop-Viewer` header.

For more information, see [Caching Content Based on Request Headers \(p. 246\)](#).

Remove Accept-Encoding Header When Compression is Not Needed

If compression is not enabled—because the origin doesn't support it, CloudFront doesn't support it, or the content is not compressible—you can increase the cache hit ratio by associating a cache behavior in your distribution to an origin that sets the Custom Origin Header as follows:

- **Header name:** `Accept-Encoding`
- **Header value:** (Keep blank)

When you use this configuration, CloudFront removes the `Accept-Encoding` header from the cache key and doesn't include the header in origin requests. This configuration applies to all content that CloudFront serves with the distribution from that origin.

Serving Media Content by Using HTTP

For information about optimizing video on demand (VOD) and streaming video content, see [Video on Demand and Live Streaming Video with CloudFront \(p. 297\)](#).

Optimizing High Availability with CloudFront Origin Failover

You can set up CloudFront with origin failover for scenarios that require high availability. To get started, you create an *origin group* with two origins: a primary and a secondary. If the primary origin is unavailable, or returns specific HTTP response status codes that indicate a failure, CloudFront automatically switches to the secondary origin.

To set up origin failover, you must have a distribution with at least two origins. Next, you create an origin group for your distribution that includes two origins, setting one as the primary. Finally, you create or update a cache behavior to use the origin group.

To see the steps for setting up origin groups and configuring specific origin failover options, see [Creating an Origin Group \(p. 232\)](#).

After you configure origin failover for a cache behavior, CloudFront does the following for viewer requests:

- When there's a cache hit, CloudFront returns the requested file.
- When there's a cache miss, CloudFront routes the request to the primary origin in the origin group.
- When the primary origin returns a status code that is not configured for failover, such as an HTTP 2xx or 3xx status code, CloudFront serves the requested content to the viewer.
- When any of the following occur:
 - The primary origin returns an HTTP status code that you've configured for failover
 - CloudFront fails to connect to the primary origin
 - The response from the primary origin takes too long (times out)

Then CloudFront routes the request to the secondary origin in the origin group.

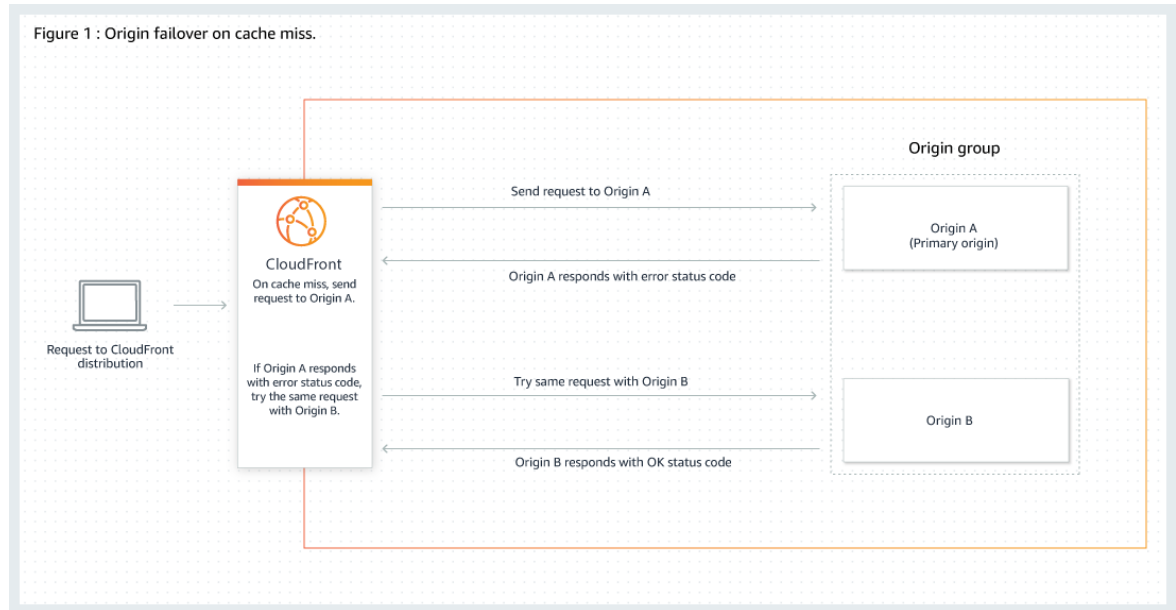
Note

For some use cases, like streaming video content, you might want CloudFront to fail over to the secondary origin quickly. To adjust how quickly CloudFront fails over to the secondary origin, see [Controlling Origin Timeouts and Attempts \(p. 233\)](#).

CloudFront routes all incoming requests to the primary origin, even when a previous request failed over to the secondary origin. CloudFront only sends requests to the secondary origin after a request to the primary origin fails.

CloudFront fails over to the secondary origin only when the HTTP method of the viewer request is `GET`, `HEAD`, or `OPTIONS`. CloudFront does not fail over when the viewer sends a different HTTP method (for example `POST`, `PUT`, and so on).

The following diagram illustrates how origin failover works.



For more information, see the following:

- **Cache hits and misses:** [How Caching Works with CloudFront Edge Caches \(p. 228\)](#)
- **Request and response behavior with origin failover:** [Request and Response Behavior for Origin Groups \(p. 282\)](#)

Topics

- [Creating an Origin Group \(p. 232\)](#)
- [Controlling Origin Timeouts and Attempts \(p. 233\)](#)
- [Use Origin Failover with Lambda@Edge Functions \(p. 234\)](#)
- [Use Custom Error Pages with Origin Failover \(p. 236\)](#)

Creating an Origin Group

To create an origin group

1. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
2. Choose the distribution that you want to create the origin group for.
3. Choose the **Origins and Origin Groups** tab.
4. On the **Origin and Origin Groups** tab, choose **Create Origin Group**.
5. Choose the origins for the origin group. After you add origins, use the arrows to set the priority—that is, which origin is primary and which is secondary.

6. Choose the HTTP status codes to use as failover criteria. You can choose any combination of the following status codes: 500, 502, 503, 504, 404, or 403. When CloudFront receives a response with one of the status codes that you specify, it fails over to the secondary origin.

Note

CloudFront fails over to the secondary origin only when the HTTP method of the viewer request is `GET`, `HEAD`, or `OPTIONS`. CloudFront does not fail over when the viewer sends a different HTTP method (for example `POST`, `PUT`, and so on).

7. Enter a unique identifier for the origin group. You can't use an identifier that's already in use for a different origin or origin group in your AWS account.

For information about specifying an origin group for a distribution, see [Origin ID \(p. 42\)](#).

Controlling Origin Timeouts and Attempts

By default, CloudFront tries to connect to the primary origin in an origin group for as long as 30 seconds (3 connection attempts of 10 seconds each) before failing over to the secondary origin. For some use cases, like streaming video content, you might want CloudFront to fail over to the secondary origin more quickly. You can adjust the following settings to affect how quickly CloudFront fails over to the secondary origin. If the origin is a secondary origin, or an origin that is not part of an origin group, these settings affect how quickly CloudFront returns an HTTP 504 response to the viewer.

To fail over more quickly, specify a shorter connection timeout, fewer connection attempts, or both. For custom origins (including Amazon S3 bucket origins that *are* configured with static website hosting), you can also adjust the origin response timeout.

Origin connection timeout

The origin connection timeout setting affects how long CloudFront waits when trying to establish a connection to the origin. By default, CloudFront waits 10 seconds to establish a connection, but you can specify 1–10 seconds (inclusive). For more information, see [Origin Connection Timeout \(p. 43\)](#).

Origin connection attempts

The origin connection attempts setting affects the number of times that CloudFront attempts to connect to the origin. By default, CloudFront tries 3 times to connect, but you can specify 1–3 (inclusive). For more information, see [Origin Connection Attempts \(p. 42\)](#).

For a custom origin (including an Amazon S3 bucket that's configured with static website hosting), this setting also affects the number of times that CloudFront attempts to get a response from the origin in the case of an origin response timeout.

Origin response timeout

Note

This applies only to custom origins.

The origin response timeout setting affects how long CloudFront waits to receive a response (or to receive the complete response) from the origin. By default, CloudFront waits for 30 seconds, but you can specify 1–60 seconds (inclusive). For more information, see [Origin Response Timeout \(p. 45\)](#).

How to change these settings

To change these settings in the [CloudFront console](#)

- For a new origin or a new distribution, you specify these values when you create the resource.
- For an existing origin in an existing distribution, you specify these values when you edit the origin.

For more information, see [Values That You Specify When You Create or Update a Distribution \(p. 38\)](#).


Use Origin Failover with Lambda@Edge Functions

You can use Lambda@Edge functions with CloudFront distributions that you've set up with origin groups. To use a Lambda function, specify it in an [origin request or origin response trigger \(p. 339\)](#) for an origin group when you create the cache behavior. When you use a Lambda@Edge function with an origin group, the function can be triggered twice for a single viewer request. For example, consider this scenario:

1. You create a Lambda@Edge function with an origin request trigger.
2. The Lambda function is triggered once when CloudFront sends a request to the primary origin (on a cache miss).
3. The primary origin responds with an HTTP status code that's configured for failover.
4. The Lambda function is triggered again when CloudFront sends the same request to the secondary origin.

The following diagram illustrates how origin failover works when you include a Lambda@Edge function in an origin request or response trigger.

Figure 2 : Origin failover with Lambda@Edge functions triggered on c



For more information about using Lambda@Edge triggers, see [Adding Triggers for a Lambda@Edge Function](#) (p. 339).

Use Custom Error Pages with Origin Failover

You can use custom error pages with origin groups similarly to how you use them with origins that are not set up for origin failover.

When you use origin failover, you can configure CloudFront to return a custom error page for the primary or secondary origin (or both):

- **Return a custom error page for the primary origin** – If the primary origin returns an HTTP status code that's not configured for failover, CloudFront returns the custom error page to viewers.
- **Return a custom error page for the secondary origin** – If CloudFront receives a failure status code from the secondary origin, CloudFront returns the custom error page.

For more information about using custom error pages with CloudFront, see [Generating Custom Error Responses](#) (p. 291).

Managing How Long Content Stays in an Edge Cache (Expiration)

You can control how long your files stay in a CloudFront cache before CloudFront forwards another request to your origin. Reducing the duration allows you to serve dynamic content. Increasing the duration means that your users get better performance because your files are more likely to be served directly from the edge cache. A longer duration also reduces the load on your origin.

Typically, CloudFront serves a file from an edge location until the cache duration that you specified passes—that is, until the file expires. After it expires, the next time the edge location gets a user request for the file, CloudFront forwards the request to the origin server to verify that the cache contains the latest version of the file. The response from the origin depends on whether the file has changed:

- If the CloudFront cache already has the latest version, the origin returns a status code 304 `Not Modified`.
- If the CloudFront cache does not have the latest version, the origin returns a status code 200 `OK` and the latest version of the file.

If a file in an edge location isn't frequently requested, CloudFront might evict the file—remove the file before its expiration date—to make room for files that have been requested more recently.

By default, each file automatically expires after 24 hours, but you can change the default behavior in two ways:

- To change the cache duration for all files that match the same path pattern, you can change the CloudFront settings for **Minimum TTL**, **Maximum TTL**, and **Default TTL** for a cache behavior. For information about the individual settings, see [Minimum TTL](#), [Maximum TTL](#), and [Default TTL in Values That You Specify When You Create or Update a Distribution](#) (p. 38). To use these settings, you must choose the **Customize** option for the **Object Caching** setting when you create or update your CloudFront distribution. For more information, see [Object Caching in Values That You Specify When You Create or Update a Distribution](#) (p. 38).
- To change the cache duration for an individual file, you can configure your origin to add a `Cache-Control max-age` or `Cache-Control s-maxage` directive, or an `Expires` header field to the file. For more information, see [Using Headers to Control Cache Duration for Individual Objects](#) (p. 237).

For more information about how **Minimum TTL**, **Default TTL**, and **Maximum TTL** interact with `Cache-Control max-age` and `Cache-Control s-maxage` directives and the `Expires` header field, see [Specifying the Amount of Time that CloudFront Caches Objects for Web Distributions \(p. 238\)](#).

You can also control how long errors (for example, `404`, `Not Found`) stay in a CloudFront cache before CloudFront tries again to get the requested object by forwarding another request to your origin. For more information, see [How CloudFront Processes and Caches HTTP 4xx and 5xx Status Codes from Your Origin \(p. 286\)](#).

Topics

- [Using Headers to Control Cache Duration for Individual Objects \(p. 237\)](#)
- [Specifying the Amount of Time that CloudFront Caches Objects for Web Distributions \(p. 238\)](#)
- [Specifying the Minimum Time that CloudFront Caches Objects for RTMP Distributions \(p. 240\)](#)
- [Adding Headers to Your Objects Using the Amazon S3 Console \(p. 241\)](#)

Using Headers to Control Cache Duration for Individual Objects

You can use the `Cache-Control` and `Expires` headers to control how long objects stay in the cache. Settings for **Minimum TTL**, **Default TTL**, and **Maximum TTL** also affect cache duration, but here's an overview of how headers can affect cache duration:

- The `Cache-Control max-age` directive lets you specify how long (in seconds) that you want an object to remain in the cache before CloudFront gets the object again from the origin server. The minimum expiration time CloudFront supports is 0 seconds for web distributions and 3600 seconds for RTMP distributions. The maximum value is 100 years. Specify the value in the following format:

```
Cache-Control: max-age=seconds
```

For example, the following directive tells CloudFront to keep the associated object in the cache for 3600 seconds (one hour):

```
Cache-Control: max-age=3600
```

If you want objects to stay in CloudFront edge caches for a different duration than they stay in browser caches, you can use the `Cache-Control max-age` and `Cache-Control s-maxage` directives together. For more information, see [Specifying the Amount of Time that CloudFront Caches Objects for Web Distributions \(p. 238\)](#).

- The `Expires` header field lets you specify an expiration date and time using the format specified in [RFC 2616, Hypertext Transfer Protocol -- HTTP/1.1 Section 3.3.1, Full Date](#), for example:

```
Sat, 27 Jun 2015 23:59:59 GMT
```

We recommend that you use the `Cache-Control max-age` directive instead of the `Expires` header field to control object caching. If you specify values both for `Cache-Control max-age` and for `Expires`, CloudFront uses only the value of `Cache-Control max-age`.

For more information, see [Specifying the Amount of Time that CloudFront Caches Objects for Web Distributions \(p. 238\)](#).

You cannot use the HTTP `Cache-Control` or `Pragma` header fields in a `GET` request from a viewer to force CloudFront to go back to the origin server for the object. CloudFront ignores those header fields in viewer requests.

For more information about the `Cache-Control` and `Expires` header fields, see the following sections in *RFC 2616, Hypertext Transfer Protocol -- HTTP/1.1*:

- [Section 14.9 Cache Control](#)
- [Section 14.21 Expires](#)

For an example of how to add `Cache-Control` and `Expires` header fields using the AWS SDK for PHP, see [Upload an Object Using the AWS SDK for PHP](#) in the *Amazon Simple Storage Service Developer Guide*. Some third-party tools are also able to add these fields.

Specifying the Amount of Time that CloudFront Caches Objects for Web Distributions

For web distributions, you can use `Cache-Control` or `Expires` headers, and CloudFront minimum, maximum, and default TTL values to control the amount of time in seconds that CloudFront keeps an object in the cache before forwarding another request to the origin. Header values also determine how long a browser keeps an object in the cache before forwarding another request to CloudFront.

Important

If you configure CloudFront to forward all headers to your origin for a cache behavior, CloudFront never caches the associated objects. Instead, CloudFront forwards all requests for those objects to the origin. In that configuration, the value of minimum TTL must be 0. For more information, see [Caching Content Based on Request Headers](#) (p. 246).

To specify values for [Minimum TTL](#), [Maximum TTL](#), and [Default TTL](#), you must choose the **Customize** option for the [Object Caching](#) setting.

Origin Configuration	Minimum TTL = 0 Seconds	Minimum TTL > 0 Seconds
The origin adds a <code>Cache-Control max-age</code> directive to objects	CloudFront caching CloudFront caches objects for the lesser of the value of the <code>Cache-Control max-age</code> directive or the value of the CloudFront maximum TTL.	CloudFront caching CloudFront caching depends on the values of the CloudFront minimum TTL and maximum TTL and the <code>Cache-Control max-age</code> directive: <ul style="list-style-type: none">• Minimum TTL < max-age < maximum TTL CloudFront caches objects for the value of the <code>Cache-Control max-age</code> directive.• max-age < minimum TTL CloudFront caches objects for the value of the CloudFront minimum TTL.• max-age > maximum TTL CloudFront caches objects for the value of the CloudFront maximum TTL.
	Browser caching Browsers cache objects for the value of the <code>Cache-Control max-age</code> directive.	Browser caching

Origin Configuration	Minimum TTL = 0 Seconds	Minimum TTL > 0 Seconds
		Browsers cache objects for the value of the <code>Cache-Control max-age</code> directive.
The origin does not add a <code>Cache-Control max-age</code> directive to objects	<p>CloudFront caching</p> <p>CloudFront caches objects for the value of the CloudFront default TTL.</p> <p>Browser caching</p> <p>Depends on the browser.</p>	<p>CloudFront caching</p> <p>CloudFront caches objects for the greater of the value of the CloudFront minimum TTL or default TTL.</p> <p>Browser caching</p> <p>Depends on the browser.</p>
The origin adds <code>Cache-Control max-age</code> and <code>Cache-Control s-maxage</code> directives to objects	<p>CloudFront caching</p> <p>CloudFront caches objects for the lesser of the value of the <code>Cache-Control s-maxage</code> directive or the value of the CloudFront maximum TTL.</p> <p>Browser caching</p> <p>Browsers cache objects for the value of the <code>Cache-Control max-age</code> directive.</p>	<p>CloudFront caching</p> <p>CloudFront caching depends on the values of the CloudFront minimum TTL and maximum TTL and the <code>Cache-Control s-maxage</code> directive:</p> <ul style="list-style-type: none"> • Minimum TTL < s-maxage < maximum TTL <p>CloudFront caches objects for the value of the <code>Cache-Control s-maxage</code> directive.</p> <ul style="list-style-type: none"> • s-maxage < minimum TTL <p>CloudFront caches objects for the value of the CloudFront minimum TTL.</p> <ul style="list-style-type: none"> • s-maxage > maximum TTL <p>CloudFront caches objects for the value of the CloudFront maximum TTL.</p> <p>Browser caching</p> <p>Browsers cache objects for the value of the <code>Cache-Control max-age</code> directive.</p>

Origin Configuration	Minimum TTL = 0 Seconds	Minimum TTL > 0 Seconds
The origin adds an <code>Expires</code> header to objects	<p>CloudFront caching</p> <p>CloudFront caches objects until the date in the <code>Expires</code> header or for the value of the CloudFront maximum TTL, whichever is sooner.</p> <p>Browser caching</p> <p>Browsers cache objects until the date in the <code>Expires</code> header.</p>	<p>CloudFront caching</p> <p>CloudFront caching depends on the values of the CloudFront minimum TTL and maximum TTL and the <code>Expires</code> header:</p> <ul style="list-style-type: none"> • Minimum TTL < <code>Expires</code> < maximum TTL <p>CloudFront caches objects until the date and time in the <code>Expires</code> header.</p> <ul style="list-style-type: none"> • <code>Expires</code> < minimum TTL <p>CloudFront caches objects for the value of the CloudFront minimum TTL.</p> <ul style="list-style-type: none"> • <code>Expires</code> > maximum TTL <p>CloudFront caches objects for the value of the CloudFront maximum TTL.</p> <p>Browser caching</p> <p>Browsers cache objects until the date and time in the <code>Expires</code> header.</p>
Origin adds <code>Cache-Control: no-cache, no-store, and/or private</code> directives to objects	<p>CloudFront and browsers respect the headers.</p> <p>For an exception to how CloudFront handles the <code>Cache-Control: no-cache</code> header, see Simultaneous Requests for the Same Object (Traffic Spikes) (p. 278).</p>	<p>CloudFront caching</p> <p>CloudFront caches objects for the value of the CloudFront minimum TTL.</p> <p>Browser caching</p> <p>Browsers respect the headers.</p>

For information about how to change settings for web distributions using the CloudFront console, see [Updating a Distribution](#) (p. 63). For information about how to change settings for web distributions using the CloudFront API, see [PUT Config](#).

Specifying the Minimum Time that CloudFront Caches Objects for RTMP Distributions

For RTMP distributions, CloudFront keeps objects in edge caches for 24 hours by default. You can add `Cache-Control` or `Expires` headers to your objects to change the amount of time that CloudFront keeps objects in edge caches before it forwards another request to the origin. The minimum duration is 3600 seconds (one hour). If you specify a lower value, CloudFront uses 3600 seconds.

Adding Headers to Your Objects Using the Amazon S3 Console

Note

On the Amazon S3 console, you can only add headers to one object at a time, but with some third-party tools, you can add headers to multiple Amazon S3 objects at a time. For more information about third-party tools that support Amazon S3, search on the web for **AWS S3 third party tools**.

To add a **Cache-Control** or **Expires** header field to Amazon S3 objects using the Amazon S3 console

1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3>.
2. In the Amazon S3 console, in the **Bucket name** list, choose the name of the bucket that contains the files.
3. In the **Name** list, choose the name of the object that you want to add a header to.
4. Choose **Properties**, and then choose **Metadata**.
5. Choose **Add Metadata**, and then in the **Key** menu, choose **Cache-Control** or **Expires**.
6. In the **Value** field, type one of the following:
 - For a **Cache-Control** field, type:

`max-age=number of seconds that you want objects to stay in a CloudFront edge cache`
 - For an **Expires** field, type a date and time in HTML format.
7. Choose **Save**.

Caching Content Based on Query String Parameters

Some web applications use query strings to send information to the origin. A query string is the part of a web request that appears after a `?` character; the string can contain one or more parameters, separated by `&` characters. In the following example, the query string includes two parameters, `color=red` and `size=large`:

```
http://d1111111abcdef8.cloudfront.net/images/image.jpg?color=red&size=large
```

For web distributions, you can choose if you want CloudFront to forward query strings to your origin and whether to cache your content based on all parameters or on selected parameters. Why might this be useful? Consider the following example.

Suppose that your website is available in five languages. The directory structure and file names for all five versions of the website are identical. As a user views your website, requests that are forwarded to CloudFront include a language query string parameter based on the language that the user chose. You can configure CloudFront to forward query strings to the origin and to cache based on the language parameter. If you configure your web server to return the version of a given page that corresponds with the selected language, CloudFront caches each language version separately, based on the value of the language query string parameter.

In this example, if the main page for your website is `main.html`, the following five requests cause CloudFront to cache `main.html` five times, once for each value of the language query string parameter:

- `http://d1111111abcdef8.cloudfront.net/main.html?language=de`
- `http://d1111111abcdef8.cloudfront.net/main.html?language=en`
- `http://d1111111abcdef8.cloudfront.net/main.html?language=es`
- `http://d1111111abcdef8.cloudfront.net/main.html?language=fr`
- `http://d1111111abcdef8.cloudfront.net/main.html?language=jp`

Note the following:

- For RTMP distributions, you cannot configure CloudFront to forward query string parameters to your origin. If you have an RTMP distribution, before CloudFront forwards a request to the origin server, it removes any query string parameters.
- Some HTTP servers don't process query string parameters and, therefore, don't return different versions of an object based on parameter values. For these origins, if you configure CloudFront to forward query string parameters to the origin, CloudFront still caches based on the parameter values even though the origin returns identical versions of the object to CloudFront for every parameter value.
- For query string parameters to work as described in the example above with the languages, you must use the & character as the delimiter between query string parameters. If you use a different delimiter, you may get unexpected results, depending on which parameters you specify for CloudFront to use as a basis for caching, and the order in which the parameters appear in the query string.

The following examples show what happens if you use a different delimiter and you configure CloudFront to cache based only on the `color` parameter:

- In the following request, CloudFront caches your content based on the value of the `color` parameter, but CloudFront interprets the value as `red;size=large`:

```
http://d1111111abcdef8.cloudfront.net/images/image.jpg?color=red;size=large
```

- In the following request, CloudFront caches your content but doesn't base caching on the query string parameters. This is because you configured CloudFront to cache based on the `color` parameter, but CloudFront interprets the following string as containing only a `size` parameter that has a value of `large;color=red`:

```
http://d1111111abcdef8.cloudfront.net/images/image.jpg?size=large;color=red
```

You can configure CloudFront do one of the following:

- Don't forward query strings to the origin at all. If you don't forward query strings, CloudFront doesn't cache based on query string parameters.
- Forward query strings to the origin, and cache based on all parameters in the query string.
- Forward query strings to the origin, and cache based on specified parameters in the query string.

For more information, see [Optimizing Caching \(p. 243\)](#).

Topics

- [Console and API Settings for Query String Forwarding and Caching \(p. 243\)](#)
- [Optimizing Caching \(p. 243\)](#)
- [Query String Parameters and CloudFront Access Logs \(p. 244\)](#)

Console and API Settings for Query String Forwarding and Caching

To configure query string forwarding and caching in the CloudFront console, see the following settings in [Values That You Specify When You Create or Update a Distribution](#) (p. 38):

- [Query String Forwarding and Caching](#) (p. 52)
- [Query String Whitelist](#) (p. 52)

To configure query string forwarding and caching with the CloudFront API, see the following settings in [DistributionConfig Complex Type](#) and in [DistributionConfigWithTags Complex Type](#) in the *Amazon CloudFront API Reference*:

- `QueryString`
- `QueryStringCacheKeys`

Optimizing Caching

When you configure CloudFront to cache based on query string parameters, you can take the following steps to reduce the number of requests that CloudFront forwards to your origin. When CloudFront edge locations serve objects, you reduce the load on your origin server and reduce latency because objects are served from locations that are closer to your users.

Cache Based Only on Parameters for Which Your Origin Returns Different Versions of an Object

For each query string parameter that your web application forwards to CloudFront, CloudFront forwards requests to your origin for every parameter value and caches a separate version of the object for every parameter value. This is true even if your origin always returns the same object regardless of the parameter value. For multiple parameters, the number of requests and the number of objects multiply. For example, if requests for an object include two parameters that each have three different values, CloudFront caches six versions of that object, assuming you follow the other recommendations in this section.

We recommend that you configure CloudFront to cache based only on the query string parameters for which your origin returns different versions, and that you carefully consider the merits of caching based on each parameter. For example, suppose you have a retail website. You have pictures of a jacket in six different colors, and the jacket comes in 10 different sizes. The pictures that you have of the jacket show the different colors but not the different sizes. To optimize caching, you should configure CloudFront to cache based only on the color parameter, not on the size parameter. This increases the likelihood that CloudFront can serve a request from the cache, which improves performance and reduces the load on your origin.

Always List Parameters in the Same Order

The order of parameters matters in query strings. In the following example, the query strings are identical except that the parameters are in a different order. This causes CloudFront to forward two separate requests for `image.jpg` to your origin and to cache two separate versions of the object:

- `http://d111111abcdef8.cloudfront.net/images/image.jpg?color=red&size=large`
- `http://d111111abcdef8.cloudfront.net/images/image.jpg?size=large&color=red`

We recommend that you always list parameter names in the same order, such as alphabetical order.

Always Use the Same Case for Parameter Names and Values

CloudFront considers the case of parameter names and values when caching based on query string parameters. In the following example, the query strings are identical except for the case

of parameter names and values. This causes CloudFront to forward four separate requests for `image.jpg` to your origin and to cache four separate versions of the object:

- `http://d111111abcdef8.cloudfront.net/images/image.jpg?color=red`
- `http://d111111abcdef8.cloudfront.net/images/image.jpg?color=Red`
- `http://d111111abcdef8.cloudfront.net/images/image.jpg?Color=red`
- `http://d111111abcdef8.cloudfront.net/images/image.jpg?Color=Red`

We recommend that you use case consistently for parameter names and values, such as all lowercase.

Don't Use Parameter Names that Conflict with Signed URLs

If you're using signed URLs to restrict access to your content (if you added trusted signers to your distribution), CloudFront removes the following query string parameters before forwarding the rest of the URL to your origin:

- `Expires`
- `Key-Pair-Id`
- `Policy`
- `Signature`

If you're using signed URLs and you want to configure CloudFront to forward query strings to your origin, your own query string parameters cannot be named `Expires`, `Key-Pair-Id`, `Policy`, or `Signature`.

Query String Parameters and CloudFront Access Logs

For web and RTMP distributions, if you enable logging, CloudFront logs the full URL, including query string parameters. For web distributions, this is true regardless of whether you have configured CloudFront to forward query strings to the origin. For more information about CloudFront logging, see [Configuring and Using Standard Logs \(Access Logs\)](#) (p. 439).

Caching Content Based on Cookies

By default, CloudFront doesn't consider cookies when processing requests and responses, or when caching your objects in edge locations. If CloudFront receives two requests that are identical except for what's in the `Cookie` header, then, by default, CloudFront treats the requests as identical and returns the same object for both requests.

You can configure CloudFront to forward to your origin some or all of the cookies in viewer requests, and to cache separate versions of your objects based on the cookie values that it forwards. When you do this, CloudFront uses some or all of the cookies in viewer requests—whichever ones it's configured to forward—to uniquely identify an object in the cache. (For RTMP distributions, you cannot configure CloudFront to process cookies and CloudFront does not cache cookies in edge caches.)

For example, suppose that requests for `locations.html` contain a `country` cookie that has a value of either `uk` or `fr`. When you configure CloudFront to cache your objects based on the value of the `country` cookie, CloudFront forwards requests for `locations.html` to the origin and includes the `country` cookie and its value. Your origin returns `locations.html`, and CloudFront caches the object once for requests in which the value of the `country` cookie is `uk` and once for requests in which the value is `fr`.

Important

Amazon S3 and some HTTP servers don't process cookies. Don't configure CloudFront to forward cookies to an origin that doesn't process cookies or doesn't vary its response based on cookies. That can cause CloudFront to forward more requests to the origin for the same object,

which slows performance and increases the load on the origin. If, considering the previous example, your origin doesn't process the `country` cookie or always returns the same version of `locations.html` to CloudFront regardless of the value of the `country` cookie, don't configure CloudFront to forward that cookie.

Conversely, if your custom origin depends on a particular cookie or sends different responses based on a cookie, make sure you configure CloudFront to forward that cookie to the origin. Otherwise, CloudFront removes the cookie before forwarding the request to your origin.

To configure cookie forwarding, you update your distribution's cache behavior. For more information about cache behaviors, see [Cache Behavior Settings \(p. 46\)](#), particularly the [Forward Cookies \(p. 51\)](#) and [Whitelist Cookies \(p. 51\)](#) sections.

You can configure each cache behavior to do one of the following:

- **Forward all cookies to your origin** – CloudFront includes all cookies sent by the viewer when it forwards requests to the origin. When your origin returns a response, CloudFront caches the response using the cookie names and values in the viewer request. If the origin response includes `Set-Cookie` headers, CloudFront returns them to the viewer with the requested object. CloudFront also caches the `Set-Cookie` headers with the object returned from the origin, and sends those `Set-Cookie` headers to viewers on all cache hits.
- **Forward a whitelist of cookies that you specify** – CloudFront removes any cookies that the viewer sends that aren't on the whitelist before it forwards a request to the origin. CloudFront caches the response using the whitelisted cookies names and values in the viewer request. If the origin response includes `Set-Cookie` headers, CloudFront returns them to the viewer with the requested object. CloudFront also caches the `Set-Cookie` headers with the object returned from the origin, and sends those `Set-Cookie` headers to viewers on all cache hits.

For information about specifying wildcards in cookie names in a whitelist, see [Whitelist Cookies \(p. 51\)](#).

For the current quota on the number of cookie names that you can whitelist for each cache behavior, or to request a higher quota, see [Quotas on Whitelisted Query Strings \(Web Distributions Only\) \(p. 499\)](#).

- **Don't forward cookies to your origin** – CloudFront doesn't cache your objects based on cookie sent by the viewer. In addition, CloudFront removes cookies before forwarding requests to your origin, and removes `Set-Cookie` headers from responses before returning responses to your viewers.

Note the following about specifying the cookies that you want to forward:

Access logs

If you configure CloudFront to log requests and to log cookies, CloudFront logs all cookies and all cookie attributes, even if you configure CloudFront not to forward cookies to your origin or if you configure CloudFront to forward only a whitelist of specific cookies. For more information about CloudFront logging, see [Configuring and Using Standard Logs \(Access Logs\) \(p. 439\)](#).

Case sensitivity

Cookie names and values are both case-sensitive. For example, if CloudFront is configured to forward all cookies, and two viewer requests for the same object have cookies that are identical except for case, CloudFront caches the object twice.

CloudFront sorts cookies

If CloudFront is configured to forward cookies (all or a whitelist), CloudFront sorts the cookies in natural order by cookie name before forwarding the request to your origin.

If-Modified-Since and If-None-Match

`If-Modified-Since` and `If-None-Match` conditional requests are not supported when CloudFront is configured to forward cookies (all or a whitelist).

Standard name–value pair format is required

CloudFront forwards a cookie header only if the value conforms to the [standard name–value pair format](#), for example: "Cookie: cookie1=value1; cookie2=value2"

Disable caching of Set–Cookie headers

If CloudFront is configured to forward cookies to the origin (all or a whitelist), it also caches the Set–Cookie headers received in the origin response. CloudFront includes these Set–Cookie headers in its response to the original viewer, and also includes them in subsequent responses that are served from the CloudFront cache.

If you want to receive cookies at your origin but you don't want CloudFront to cache the Set–Cookie headers in your origin's responses, configure your origin to add a Cache–Control header with a no–cache directive that specifies Set–Cookie as a field name. For example: Cache–Control: no–cache="Set–Cookie". For more information, see [Response Cache–Control Directives](#) in the *Hypertext Transfer Protocol (HTTP/1.1): Caching* standard.

Maximum length of cookie names

If you configure CloudFront to forward a whitelist of specific cookies to your origin, the total number of bytes in all of the cookie names that you configure CloudFront to forward can't exceed 512 minus the number of cookies that you're forwarding. For example, if you configure CloudFront to forward 10 cookies to your origin, the combined length of the names of the 10 cookies can't exceed 502 bytes (512 – 10).

If you configure CloudFront to forward all cookies to your origin, the length of cookie names doesn't matter.

For information about using the CloudFront console to update a distribution so CloudFront forwards cookies to the origin, see [Updating a Distribution](#) (p. 63). For information about using the CloudFront API to update a distribution, see [UpdateDistribution](#) in the *Amazon CloudFront API Reference*.

Caching Content Based on Request Headers

For web distributions, CloudFront lets you choose whether you want CloudFront to forward headers to your origin and to cache separate versions of a specified object based on the header values in viewer requests. This allows you to serve different versions of your content based on the device the user is using, the location of the viewer, the language the viewer is using, and a variety of other criteria. For RTMP distributions, you cannot configure CloudFront to cache based on header values.

Note

For RTMP distributions, you cannot configure CloudFront to cache your content based on the headers in viewer requests.

Topics

- [Headers and Distributions - Overview](#) (p. 247)
- [Selecting the Headers to Base Caching On](#) (p. 247)
- [Configuring CloudFront to Respect CORS Settings](#) (p. 248)
- [Configuring Caching Based on the Device Type](#) (p. 248)
- [Configuring Caching Based on the Language of the Viewer](#) (p. 249)
- [Configuring Caching Based on the Location of the Viewer](#) (p. 249)
- [Configuring Caching Based on the Protocol of the Request](#) (p. 249)
- [Configuring Caching For Compressed Files](#) (p. 249)
- [How Caching Based on Headers Affects Performance](#) (p. 249)

- [How the Case of Headers and Header Values Affects Caching \(p. 250\)](#)
- [Headers that CloudFront Returns to the Viewer \(p. 250\)](#)

Headers and Distributions - Overview

By default, CloudFront doesn't consider headers when caching your objects in edge locations. If your origin returns two objects and they differ only by the values in the request headers, CloudFront caches only one version of the object.

You can configure CloudFront to forward headers to the origin, which causes CloudFront to cache multiple versions of an object based on the values in one or more request headers. To configure CloudFront to cache objects based on the values of specific headers, you specify cache behavior settings for your distribution. For more information, see [Cache Based on Selected Request Headers](#).

For example, suppose viewer requests for `logo.jpg` contain a custom `Product` header that has a value of either `Acme` or `Apex`. When you configure CloudFront to cache your objects based on the value of the `Product` header, CloudFront forwards requests for `logo.jpg` to the origin and includes the `Product` header and header values. CloudFront caches `logo.jpg` once for requests in which the value of the `Product` header is `Acme` and once for requests in which the value is `Apex`.

You can configure each cache behavior in a web distribution to do one of the following:

- Forward all headers to your origin

Important

If you configure CloudFront to forward all headers to your origin, CloudFront doesn't cache the objects associated with this cache behavior. Instead, it sends every request to the origin.

- Forward a whitelist of headers that you specify. CloudFront caches your objects based on the values in all of the specified headers. CloudFront also forwards the headers that it forwards by default, but it caches your objects based only on the headers that you specify.
- Forward only the default headers. In this configuration, CloudFront doesn't cache your objects based on the values in the request headers.

For the current quota on the number of headers that you can whitelist for each cache behavior or to request a higher quota, see [Quotas on Custom Headers \(Web Distributions Only\) \(p. 499\)](#).

For information about using the CloudFront console to update a distribution so CloudFront forwards headers to the origin, see [Updating a Distribution \(p. 63\)](#). For information about using the CloudFront API to update an existing distribution, see [Update Distribution](#) in the *Amazon CloudFront API Reference*.

Selecting the Headers to Base Caching On

The headers that you can forward to the origin and that CloudFront bases caching on depend on whether your origin is an Amazon S3 bucket or a custom origin.

- **Amazon S3** – You can configure CloudFront to forward and to cache your objects based on a number of specific headers (see the following list of exceptions). However, we recommend that you avoid whitelisting headers with an Amazon S3 origin unless you need to implement cross-origin resource sharing (CORS) or you want to personalize content by using `Lambda@Edge` in origin-facing events.
 - To configure CORS, you must forward headers that allow CloudFront to distribute content for websites that are enabled for cross-origin resource sharing (CORS). For more information, see [Configuring CloudFront to Respect CORS Settings \(p. 248\)](#).
 - To personalize content by using headers that you forward to your Amazon S3 origin, you write and add `Lambda@Edge` functions and associate them with your CloudFront distribution to be triggered

by an origin-facing event. For more information about working with headers to personalize content, see [Personalize Content by Country or Device Type Headers - Examples](#) (p. 377).

We recommend that you avoid whitelisting headers that you aren't using to personalize content because forwarding extra headers can reduce your cache hit ratio. That is, CloudFront can't serve as many requests from edge caches, as a proportion of all requests.

- **Custom origin** – You can configure CloudFront to cache based on the value of any request header except the following:
 - Connection
 - Cookie – If you want to forward and cache based on cookies, you use a separate setting in your distribution. For more information, see [Caching Content Based on Cookies](#) (p. 244).
 - Host (for Amazon S3 origins)
 - Proxy-Authorization
 - TE
 - Upgrade

You can configure CloudFront to cache objects based on values in the Date and User-Agent headers, but we don't recommend it. These headers have numerous possible values, and caching based on their values could cause CloudFront to forward significantly more requests to your origin.

For a full list of HTTP request headers and how CloudFront processes them, see [HTTP Request Headers and CloudFront Behavior \(Custom and S3 Origins\)](#) (p. 273).

Configuring CloudFront to Respect CORS Settings

If you have enabled cross-origin resource sharing (CORS) on an Amazon S3 bucket or a custom origin, you must choose specific headers to forward, to respect the CORS settings. The headers that you must forward differ depending on the origin (Amazon S3 or custom) and whether you want to cache OPTIONS responses.

Amazon S3

- If you want OPTIONS responses to be cached, do the following:
 - Choose the options for default cache behavior settings that enable caching for OPTIONS responses.
 - Configure CloudFront to forward the following headers: Origin, Access-Control-Request-Headers, and Access-Control-Request-Method.
- If you don't want OPTIONS responses to be cached, configure CloudFront to forward the Origin header, together with any other headers required by your origin (for example, Access-Control-Request-Headers, Access-Control-Request-Method, or others).

Custom origins – Forward the Origin header along with any other headers required by your origin.

You configure CloudFront to forward headers by whitelisting the headers in a cache behavior for your CloudFront distribution. For more information about working with header forwarding, see [Headers and Distributions - Overview](#) (p. 247).

For more information about CORS and Amazon S3, see [Enabling Cross-Origin Resource Sharing](#) in the *Amazon Simple Storage Service Developer Guide*.

Configuring Caching Based on the Device Type

If you want CloudFront to cache different versions of your objects based on the device a user is using to view your content, configure CloudFront to forward the applicable headers to your custom origin:

- `CloudFront-Is-Desktop-Viewer`
- `CloudFront-Is-Mobile-Viewer`
- `CloudFront-Is-SmartTV-Viewer`
- `CloudFront-Is-Tablet-Viewer`

Based on the value of the `User-Agent` header, CloudFront sets the value of these headers to `true` or `false` before forwarding the request to your origin. If a device falls into more than one category, more than one value might be `true`. For example, for some tablet devices, CloudFront might set both `CloudFront-Is-Mobile-Viewer` and `CloudFront-Is-Tablet-Viewer` to `true`.

Configuring Caching Based on the Language of the Viewer

If you want CloudFront to cache different versions of your objects based on the language specified in the request, configure CloudFront to forward the `Accept-Language` header to your origin.

Configuring Caching Based on the Location of the Viewer

If you want CloudFront to cache different versions of your objects based on the country that the request came from, configure CloudFront to forward the `CloudFront-Viewer-Country` header to your origin. CloudFront automatically converts the IP address that the request came from into a two-letter country code. For an easy-to-use list of country codes, sortable by code and by country name, see the Wikipedia entry [ISO 3166-1 alpha-2](#).

Configuring Caching Based on the Protocol of the Request

If you want CloudFront to cache different versions of your objects based on the protocol of the request, HTTP or HTTPS, configure CloudFront to forward the `CloudFront-Forwarded-Proto` header to your origin.

Configuring Caching For Compressed Files

If your origin supports brotli compression, you can whitelist the `Accept-Encoding` header and cache based on the header. Configure caching based on `Accept-Encoding` only if your origin serves different content based on the header.

How Caching Based on Headers Affects Performance

When you configure CloudFront to cache based on one or more headers and the headers have more than one possible value, CloudFront forwards more requests to your origin server for the same object. This slows performance and increases the load on your origin server. If your origin server returns the same object regardless of the value of a given header, we recommend that you don't configure CloudFront to cache based on that header.

If you configure CloudFront to forward more than one header, the order of the headers in viewer requests doesn't affect caching as long as the values are the same. For example, if one request contains the headers `A:1,B:2` and another request contains `B:2,A:1`, CloudFront caches just one copy of the object.

How the Case of Headers and Header Values Affects Caching

When CloudFront caches based on header values, it doesn't consider the case of the header name, but it does consider the case of the header value:

- If viewer requests include both `Product:Acme` and `product:Acme`, CloudFront caches an object only once. The only difference between them is the case of the header name, which doesn't affect caching.
- If viewer requests include both `Product:Acme` and `Product:acme`, CloudFront caches an object twice, because the value is `Acme` in some requests and `acme` in others.

Headers that CloudFront Returns to the Viewer

Configuring CloudFront to forward and cache headers does not affect which headers CloudFront returns to the viewer. CloudFront returns all of the headers that it gets from the origin with a few exceptions. For more information, see the applicable topic:

- **Amazon S3 origins** – See [HTTP Response Headers That CloudFront Removes or Updates \(p. 268\)](#).
- **Custom origins** – See [HTTP Response Headers that CloudFront Removes or Replaces \(p. 280\)](#).

Troubleshooting

Troubleshoot common problems you might encounter when setting up Amazon CloudFront to distribute your content or when using Lambda@Edge, and find possible solutions.

Topics

- [Troubleshooting Distribution Issues \(p. 251\)](#)
- [Troubleshooting Error Responses from Your Origin \(p. 254\)](#)
- [Load Testing CloudFront \(p. 262\)](#)

Troubleshooting Distribution Issues

Use the information here to help you diagnose and fix certificate errors, access-denied issues, or other common issues that you might encounter when setting up your website or application with Amazon CloudFront distributions.

Topics

- [CloudFront Returns an InvalidViewerCertificate Error When I Try to Add an Alternate Domain Name \(p. 251\)](#)
- [I Can't View the Files in My Distribution \(p. 252\)](#)
- [Error Message: Certificate: <certificate-id> Is Being Used by CloudFront \(p. 253\)](#)

CloudFront Returns an InvalidViewerCertificate Error When I Try to Add an Alternate Domain Name

If CloudFront returns an `InvalidViewerCertificate` error when you try to add an alternate domain name (CNAME) to your distribution, review the following information to help troubleshoot the problem. This error can indicate that one of the following issues must be resolved before you can successfully add the alternate domain name.

The following errors are listed in the order in which CloudFront checks for authorization to add an alternate domain name. This can help you troubleshoot issues because based on the error that CloudFront returns, you can tell which verification checks have completed successfully.

There's no certificate attached to your distribution.

To add an alternate domain name (CNAME), you must attach a trusted, valid certificate to your distribution. Please review the requirements, obtain a valid certificate that meets them, attach it to your distribution, and then try again. For more information, see [Requirements for Using Alternate Domain Names \(p. 78\)](#).

There are too many certificates in the certificate chain for the certificate that you've attached.

You can only have up to five certificates in a certificate chain. Reduce the number of certificates in the chain, and then try again.

The certificate chain includes one or more certificates that aren't valid for the current date.

The certificate chain for a certificate that you have added has one or more certificates that aren't valid, either because a certificate isn't valid yet or a certificate has expired. Check the **Not Valid Before** and **Not Valid After** fields in the certificates in your certificate chain to make sure that all of the certificates are valid based on the dates that you've listed.

The certificate that you've attached isn't signed by a trusted Certificate Authority (CA).

The certificate that you attach to CloudFront to verify an alternate domain name cannot be a self-signed certificate. It must be signed by a trusted CA. For more information, see [Requirements for Using Alternate Domain Names \(p. 78\)](#).

The certificate that you've attached isn't formatted correctly

The domain name and IP address format that are included in the certificate, and the format of the certificate itself, must follow the standard for certificates.

There was a CloudFront internal error.

CloudFront was blocked by an internal issue and couldn't make validation checks for certificates. In this scenario, CloudFront returns an HTTP 500 status code and indicates that there is an internal CloudFront problem with attaching the certificate. Wait a few minutes, and then try again to add the alternate domain name with the certificate.

The certificate that you've attached doesn't cover the alternate domain name that you're trying to add.

For each alternate domain name that you add, CloudFront requires that you attach a valid SSL/TLS certificate from a trusted Certificate Authority (CA) that covers the domain name, to validate your authorization to use it. Please update your certificate to include a domain name that covers the CNAME that you're trying to add. For more information and examples of using domain names with wildcards, see [Requirements for Using Alternate Domain Names \(p. 78\)](#).

I Can't View the Files in My Distribution

If you can't view the files in your CloudFront distribution, see the following topics for some common solutions.

Did You Sign Up for Both CloudFront and Amazon S3?

To use Amazon CloudFront with an Amazon S3 origin, you must sign up for both CloudFront and Amazon S3, separately. For more information about signing up for CloudFront and Amazon S3, see [Setting up Amazon CloudFront \(p. 11\)](#).

Are Your Amazon S3 Bucket and Object Permissions Set Correctly?

If you are using CloudFront with an Amazon S3 origin, the original versions of your content are stored in an S3 bucket. The easiest way to use CloudFront with Amazon S3 is to make all of your objects publicly readable in Amazon S3. To do this, you must explicitly enable public read privileges for each object that you upload to Amazon S3.

If your content is not publicly readable, you must create a CloudFront origin access identity (OAI) so that CloudFront can access it. For more information about CloudFront origin access identities, see [Restricting Access to Amazon S3 Content by Using an Origin Access Identity \(p. 209\)](#).

Object properties and bucket properties are independent. You must explicitly grant privileges to each object in Amazon S3. Objects do not inherit properties from buckets, and object properties must be set independently of the bucket.

Is Your Alternate Domain Name (CNAME) Correctly Configured?

If you already have an existing CNAME record for your domain name, update that record or replace it with a new one that points to your distribution's domain name.

Also, make sure that your CNAME record points to your distribution's domain name, not your Amazon S3 bucket. You can confirm that the CNAME record in your DNS system points to your distribution's domain name. To do so, use a DNS tool like dig. For information about dig, see <http://www.kloth.net/services/dig.php>.

The following example shows a dig request for a domain name called `images.example.com` and the relevant part of the response. Under `ANSWER SECTION`, see the line that contains `CNAME`. The CNAME record for your domain name is set up correctly if the value on the right side of `CNAME` is your CloudFront distribution's domain name. If it's your Amazon S3 origin server bucket or some other domain name, then the CNAME record is set up incorrectly.

```
[prompt]>
dig images.example.com

; <<> DiG 9.3.3rc2 <<> images.example.com
;; global options: printcmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 15917
;; flags: qr rd ra; QUERY: 1, ANSWER: 9, AUTHORITY: 2, ADDITIONAL: 0
;; QUESTION SECTION:
;images.example.com.      IN      A
;; ANSWER SECTION:
images.example.com. 10800 IN CNAME d111111abcdef8.cloudfront.net.
...
...
```

For more information about CNAMEs, see [Using Custom URLs for Files by Adding Alternate Domain Names \(CNAMEs\)](#) (p. 71).

Are You Referencing the Correct URL for Your CloudFront Distribution?

Make sure that the URL that you're referencing uses the domain name (or CNAME) of your CloudFront distribution, not your Amazon S3 bucket or custom origin.

Do You Need Help Troubleshooting a Custom Origin?

If you need AWS to help you troubleshoot a custom origin, we probably will need to inspect the `X-Amz-Cf-Id` header entries from your requests. If you are not already logging these entries, you might want to consider it for the future. For more information, see [Using Amazon EC2 or Other Custom Origins](#) (p. 68). For further help, see the [AWS Support Center](#).

I Can't View the Files in My RTMP Distribution

If you can't view the files in your RTMP distribution, are your URL and your playback client correctly configured? RTMP distributions require you to use an RTMP protocol instead of HTTP, and you must make a few minor configuration changes to your playback client. For information about creating RTMP distributions, see [Task List for Streaming Media Files Using RTMP](#) (p. 307).

Error Message: Certificate: <certificate-id> Is Being Used by CloudFront

Problem: You're trying to delete an SSL/TLS certificate from the IAM certificate store, and you're getting the message "Certificate: <certificate-id> is being used by CloudFront."

Solution: Every CloudFront web distribution must be associated either with the default CloudFront certificate or with a custom SSL/TLS certificate. Before you can delete an SSL/TLS certificate, you must either rotate the certificate (replace the current custom SSL/TLS certificate with another custom SSL/TLS certificate) or revert from using a custom SSL/TLS certificate to using the default CloudFront certificate. To fix that, complete the steps in one of the following procedures:

- [Rotating SSL/TLS Certificates \(p. 143\)](#)
- [Reverting from a Custom SSL/TLS Certificate to the Default CloudFront Certificate \(p. 144\)](#)

Troubleshooting Error Responses from Your Origin

If CloudFront requests an object from your origin, and the origin returns an HTTP 4xx or 5xx status code, there's a problem with communication between CloudFront and your origin. The following topics describe common causes for some of these HTTP status codes, and some possible solutions.

Topics

- [HTTP 400 Status Code \(Bad Request\) \(p. 254\)](#)
- [HTTP 500 Status Code \(Lambda Execution Error\) \(p. 255\)](#)
- [HTTP 502 Status Code \(Bad Gateway\) \(p. 255\)](#)
- [HTTP 502 Status Code \(Lambda Validation Error\) \(p. 258\)](#)
- [HTTP 503 Status Code \(Lambda Limit Exceeded\) \(p. 258\)](#)
- [HTTP 503 Status Code \(Service Unavailable\) \(p. 258\)](#)
- [HTTP 504 Status Code \(Gateway Timeout\) \(p. 259\)](#)

HTTP 400 Status Code (Bad Request)

Your CloudFront distribution might send error responses with HTTP status code 400 Bad Request, and a message similar to the following:

The authorization header is malformed; the region '<AWS Region>' is wrong; expecting '<AWS Region>'

For example:

The authorization header is malformed; the region 'us-east-1' is wrong; expecting 'us-west-2'

This problem can occur in the following scenario:

1. Your CloudFront distribution's origin is an Amazon S3 bucket.
2. You moved the S3 bucket from one AWS Region to another. That is, you deleted the S3 bucket, then later you created a new bucket with the same bucket name, but in a different AWS Region than where the original S3 bucket was located.

To fix this error, update your CloudFront distribution so that it finds the S3 bucket in the bucket's current AWS Region.

To update your CloudFront distribution

1. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
2. Choose the distribution that produces this error.
3. Choose **Origins and Origin Groups**.

4. Find the origin for the S3 bucket that you moved. Select the check box next to this origin, then choose **Edit**.
5. Choose **Yes, Edit**. You do not need to change any settings before choosing **Yes, Edit**.

When you complete these steps, CloudFront redeploys your distribution. The distribution's status in the CloudFront console changes to **In Progress** while the distribution is deploying. When the deployment is complete, the distribution's status changes to **Deployed**, and you should stop receiving the `AuthorizationHeaderMalformed` error responses. Even after the status changes to **Deployed**, it might take some time before you stop receiving this error.

HTTP 500 Status Code (Lambda Execution Error)

If you're using Lambda@Edge, an HTTP 500 status code can indicate that your Lambda function returned an execution error. For more information about troubleshooting Lambda@Edge errors, see [Testing and Debugging Lambda@Edge Functions](#) (p. 343).

HTTP 502 Status Code (Bad Gateway)

An HTTP 502 status code (Bad Gateway) indicates that CloudFront wasn't able to serve the requested object because it couldn't connect to the origin server.

Topics

- [SSL/TLS Negotiation Failure Between CloudFront and a Custom Origin Server](#) (p. 255)
- [Origin Is Not Responding with Supported Ciphers/Protocols](#) (p. 256)
- [SSL/TLS Certificate on the Origin Is Expired, Invalid, Self-signed, or the Certificate Chain Is in the Wrong Order](#) (p. 257)
- [Origin Is Not Responding on Specified Ports in Origin Settings](#) (p. 257)
- [CloudFront Was Not Able to Resolve Your Origin Domain Due to DNS Issues](#) (p. 257)
- [Lambda@Edge Function Validation Errors](#) (p. 258)

SSL/TLS Negotiation Failure Between CloudFront and a Custom Origin Server

If you use a custom origin and you configured CloudFront to require HTTPS between CloudFront and your origin, the problem might be mismatched domain names. The SSL/TLS certificate that is installed on your origin includes a domain name in the **Common Name** field and possibly several more in the **Subject Alternative Names** field. (CloudFront supports wildcard characters in certificate domain names.) One of the domain names in the certificate must match one or both of the following values:

- The value that you specified for **Origin Domain Name** for the applicable origin in your distribution.
- The value of the `Host` header if you configured CloudFront to forward the `Host` header to your origin. For more information about forwarding the `Host` header to your origin, see [Caching Content Based on Request Headers](#) (p. 246).

If the domain names don't match, the SSL/TLS handshake fails, and CloudFront returns an HTTP status code 502 (Bad Gateway) and sets the `X-Cache` header to `Error from cloudfront`.

To determine whether domain names in the certificate match the **Origin Domain Name** in the distribution or the `Host` header, you can use an online SSL checker or OpenSSL. If the domain names don't match, you have two options:

- Get a new SSL/TLS certificate that includes the applicable domain names.

If you use AWS Certificate Manager (ACM), see [Request a Certificate](#) in the *AWS Certificate Manager User Guide* to request a new certificate.

- Change the distribution configuration so CloudFront no longer tries to use SSL to connect with your origin.

Online SSL Checker

To find an SSL test tool, search the internet for "online ssl checker." Typically, you specify the name of your domain, and the tool returns a variety of information about your SSL/TLS certificate. Confirm that the certificate contains your domain name in the **Common Name** or **Subject Alternative Names** fields.

OpenSSL

To help troubleshoot HTTP 502 errors from CloudFront, you can use OpenSSL to try to make an SSL/TLS connection to your origin server. If OpenSSL is not able to make a connection, that can indicate a problem with your origin server's SSL/TLS configuration. If OpenSSL is able to make a connection, it returns information about the origin server's certificate, including the certificate's common name (Subject CN field) and subject alternative name (Subject Alternative Name field).

Use the following OpenSSL command to test the connection to your origin server (replace *origin domain name* with your origin server's domain name, such as example.com):

```
openssl s_client -connect origin domain name:443
```

If the following are true:

- Your origin server supports multiple domain names with multiple SSL/TLS certificates
- Your distribution is configured to forward the `Host` header to the origin

Then add the `-servername` option to the OpenSSL command, as in the following example (replace *CNAME* with the CNAME that's configured in your distribution):

```
openssl s_client -connect origin domain name:443 -servername CNAME
```

Origin Is Not Responding with Supported Ciphers/Protocols

CloudFront connects to origin servers using ciphers and protocols. For a list of the ciphers and protocols that CloudFront supports, see [Supported Ciphers and Protocols \(p. 128\)](#). If your origin does not respond with one of these ciphers or protocols in the SSL/TLS exchange, CloudFront fails to connect. You can validate that your origin supports the ciphers and protocols by using SSL Labs:

- [SSL Labs](#)

Type the domain name of your origin in the **Hostname** field, and then choose **Submit**. Review the **Common names** and **Alternative names** fields from the test to see if they match your origin's domain name.

After the test is finished, find the **Protocols** and **Cipher Suites** sections in the test results to see which ciphers or protocols are supported by your origin. Compare them with the list of [Supported Ciphers and Protocols \(p. 128\)](#).

Note

If you're using Elastic Load Balancing, see [SSL Security Policies for Elastic Load Balancing](#) in the *Elastic Load Balancing User Guide* to learn how to set the ciphers and protocols. Using the Predefined Security Policy `ELBSecurityPolicy-2016-08` gives CloudFront access to your elastic

load balancer. If you want to restrict it further using a custom policy, you must allow the ciphers that CloudFront supports.

SSL/TLS Certificate on the Origin Is Expired, Invalid, Self-signed, or the Certificate Chain Is in the Wrong Order

If the origin server returns the following, CloudFront drops the TCP connection, returns HTTP status code 502 (Bad Gateway), and sets the `X-Cache` header to `Error from cloudfront`:

- An expired certificate
- Invalid certificate
- Self-signed certificate
- Certificate chain in the wrong order

Note

If the full chain of certificates, including the intermediate certificate, is not present, CloudFront drops the TCP connection.

For information about installing an SSL/TLS certificate on your custom origin server, see [Requiring HTTPS for Communication Between CloudFront and Your Custom Origin](#) (p. 123).

Origin Is Not Responding on Specified Ports in Origin Settings

When you create an origin on your CloudFront distribution, you can set the ports that CloudFront connects to the origin with for HTTP and HTTPS traffic. By default, these are TCP 80/443. You have the option to modify these ports. If your origin is rejecting traffic on these ports for any reason, or if your backend server isn't responding on the ports, CloudFront will fail to connect.

To troubleshoot these issues, check any firewalls running in your infrastructure and validate that they are not blocking the supported IP ranges. For more information, see [AWS IP Address Ranges](#) in the *Amazon Web Services General Reference*. Additionally, verify whether your web server is running on the origin.

CloudFront Was Not Able to Resolve Your Origin Domain Due to DNS Issues

When CloudFront receives a request for an object that is expired or is not stored in its cache, it makes a request to the origin to get the updated object. To make a successful request to the origin, CloudFront performs a DNS resolution on the origin domain name. However, when the DNS service that hosts your domain is experiencing issues, CloudFront cannot resolve the domain name to get the IP address, resulting in a 502 error. To fix this issue, contact your DNS provider, or, if you are using Amazon Route 53, see [Amazon Route 53 DNS](#).

To further troubleshoot this issue, ensure that the [authoritative name servers](#) of your origin's root domain or zone apex (such as `example.com`) are functioning correctly. Your authoritative name servers then receive the request and return the IP address that is associated with the domain, and are the same as the DNS servers that you used to set up your CloudFront distribution. Use the following commands to find the name servers for your apex origin:

```
dig OriginAPEXDomainName NS +short  
nslookup -query=NS OriginAPEXDomainName
```

When you have the names of your name servers, use the following commands to query the domain name of your origin against them to make sure that each responds with an answer:

```
dig OriginDomainName @NameServerFromAbove  
nslookup OriginDomainName NameServerFromAbove
```

Lambda@Edge Function Validation Errors

If you're using Lambda@Edge, an HTTP 502 status code can indicate that your Lambda function response was incorrectly formed or included invalid content. For more information about troubleshooting Lambda@Edge errors, see [Testing and Debugging Lambda@Edge Functions](#) (p. 343).

HTTP 502 Status Code (Lambda Validation Error)

If you're using Lambda@Edge, an HTTP 502 status code can indicate that your Lambda function response was incorrectly formed or included invalid content. For more information about troubleshooting Lambda@Edge errors, see [Testing and Debugging Lambda@Edge Functions](#) (p. 343).

HTTP 503 Status Code (Lambda Limit Exceeded)

If you're using Lambda@Edge, an HTTP 503 status code can indicate that the Lambda service returned an error. The error might be caused by one of the following:

- The number of function executions exceeded one of the quotas (formerly known as limits) that Lambda sets to throttle executions in an AWS Region (concurrent executions or invocation frequency).
- The function exceeded the Lambda function timeout quota.

For more information about the AWS Lambda quotas, see [AWS Lambda Limits](#) in the *AWS Lambda Developer Guide*. For more information about troubleshooting Lambda@Edge errors, see [Testing and Debugging Lambda@Edge Functions](#) (p. 343).

HTTP 503 Status Code (Service Unavailable)

An HTTP 503 status code (Service Unavailable) typically indicates a performance issue on the origin server. In rare cases, it indicates that CloudFront temporarily can't satisfy a request because of resource constraints at an edge location.

Topics

- [Origin Server Does Not Have Enough Capacity to Support the Request Rate](#) (p. 258)
- [CloudFront Caused the Error Due to Resource Constraints at the Edge Location](#) (p. 259)

Origin Server Does Not Have Enough Capacity to Support the Request Rate

CloudFront generates this error when the origin server is overwhelmed with incoming requests. CloudFront then relays the error back to the user. To resolve this issue, try the following solutions:

- If you use Amazon S3 as your origin server, optimize the performance of Amazon S3 by following the best practices for key naming. For more information, see [Request Rate and Performance Considerations](#) in the *Amazon Simple Storage Service Developer Guide*.
- If you use Elastic Load Balancing as your origin server, see [503 Error Classic](#).
- If you use a custom origin, examine the application logs to ensure that your origin has sufficient resources, such as memory, CPU, and disk size. If you use Amazon EC2 as the backend, make sure that

the instance type has the appropriate resources to fulfill the incoming requests. For more information, see [Instance Types](#) in the *Amazon EC2 User Guide for Linux Instances*.

CloudFront Caused the Error Due to Resource Constraints at the Edge Location

You will receive this error in the rare situation that CloudFront can't route requests to the next best available edge location, and so can't satisfy a request. This error is common when you perform load testing on your CloudFront distribution. To help prevent this, follow the [Load Testing CloudFront \(p. 262\)](#) guidelines for avoiding 503 (Capacity Exceeded) errors.

If this happens in your production environment, contact [AWS Support](#).

HTTP 504 Status Code (Gateway Timeout)

An HTTP 504 status code (Gateway Timeout) indicates that when CloudFront forwarded a request to the origin (because the requested object wasn't in the edge cache), one of the following happened:

- The origin returned an HTTP 504 status code to CloudFront.
- The origin didn't respond before the request expired.

CloudFront will return an HTTP 504 status code if traffic is blocked to the origin by a firewall or security group, or if the origin isn't accessible on the internet. Check for those issues first. Then, if access isn't the problem, explore application delays and server timeouts to help you identify and fix the issues.

Topics

- [Configure the Firewall on Your Origin Server to Allow CloudFront Traffic \(p. 259\)](#)
- [Configure the Security Groups on Your Origin Server to Allow CloudFront Traffic \(p. 260\)](#)
- [Make Your Custom Origin Server Accessible on the Internet \(p. 260\)](#)
- [Find and Fix Delayed Responses from Applications on Your Origin Server \(p. 260\)](#)

Configure the Firewall on Your Origin Server to Allow CloudFront Traffic

If the firewall on your origin server blocks CloudFront traffic, CloudFront returns an HTTP 504 status code, so it's good to make sure that isn't the issue before checking for other problems.

The method that you use to determine if this is an issue with your firewall depends on what system your origin server uses:

- If you use an IPTable firewall on a Linux server, you can search for tools and information to help you work with IPTables.
- If you use Windows Firewall on a Windows server, see [Add or Edit Firewall Rule](#) on Microsoft Technet.

When you evaluate the firewall configuration on your origin server, look for any firewalls or security rules that block traffic from CloudFront edge locations, based on the [published IP address range](#).

If the CloudFront IP address range is whitelisted on your origin server, make sure to update your server's security rules to incorporate changes. You can subscribe to an Amazon Simple Notification Service (SNS) topic and receive notifications when the IP address range file is updated. After you receive

the notification, you can use code to retrieve the file, parse it, and make adjustments for your local environment." For more information, see [Subscribe to AWS Public IP Address Changes via Amazon SNS](#).

Configure the Security Groups on Your Origin Server to Allow CloudFront Traffic

If your origin uses Elastic Load Balancing, review the [ELB security groups](#) and make sure that the security groups allow inbound traffic from CloudFront.

You can also use AWS Lambda to [automatically update your security groups](#) to allow inbound traffic from CloudFront.

Make Your Custom Origin Server Accessible on the Internet

If CloudFront can't access your custom origin server because it isn't publicly available on the internet, CloudFront returns an HTTP 504 error.

CloudFront edge locations connect to origin servers through the internet. If your custom origin is on a private network, CloudFront can't reach it. Because of this, you can't use private servers, including [internal Classic Load Balancers](#), as origin servers with CloudFront.

To check that internet traffic can connect to your origin server, run the following commands (where OriginDomainName is the domain name for your server):

For HTTPS traffic:

- `nc -zv OriginDomainName 443`
- `telnet OriginDomainName 443`

For HTTP traffic:

- `nc -zv OriginDomainName 80`
- `telnet OriginDomainName 80`

Find and Fix Delayed Responses from Applications on Your Origin Server

Server timeouts are often the result of either an application taking a very long time to respond, or a timeout value that is set too low.

A quick fix to help avoid HTTP 504 errors is to simply set a higher CloudFront timeout value for your distribution. But we recommend that you first make sure that you address any performance and latency issues with the application and origin server. Then you can set a reasonable timeout value that helps prevent HTTP 504 errors and provides good responsiveness to users.

Here's an overview of the steps you can take to find performance issues and correct them:

1. Measure the typical and high-load latency (responsiveness) of your web application.
2. Add additional resources, such as CPU or memory, if needed. Take other steps to address issues, such as tuning database queries to accommodate high-load scenarios.
3. If needed, adjust the timeout value for your CloudFront web distribution.

Following are details about each step.

Measure typical and high-load latency

To determine if one or more backend web application servers are experiencing high latency, run the following Linux curl command on each server:

```
curl -w "Connect time: %{time_connect} Time to first byte:
%{time_starttransfer} Total time: %{time_total} \n" -o /dev/null https://
www.example.com/yourobject
```

Note

If you run Windows on your servers, you can search for and download curl for Windows to run a similar command.

As you measure and evaluate the latency of an application that runs on your server, keep in mind the following:

- Latency values are relative to each application. However, a Time to First Byte in milliseconds rather than seconds or more, is reasonable.
- If you measure the application latency under normal load and it's fine, be aware that viewers might still experience timeouts under high load. When there is high demand, servers can have delayed responses or not respond at all. To help prevent high-load latency issues, check your server's resources such as CPU, memory, and disk reads and writes to make sure that your servers have the capacity to scale for high load.

You can run the following Linux command to check the memory that is used by Apache processes:

```
watch -n 1 "echo -n 'Apache Processes: ' && ps -C apache2 --no-headers | wc -
l && free -m"
```

- High CPU utilization on the server can significantly reduce an application's performance. If you use an Amazon EC2 instance for your backend server, review the CloudWatch metrics for the server to check the CPU utilization. For more information, see the [Amazon CloudWatch User Guide](#). Or if you're using your own server, refer to the server Help documentation for instructions on how to check CPU utilization.
- Check for other potential issues under high loads, such as database queries that run slowly when there's a high volume of requests.

Add resources, and tune servers and databases

After you evaluate the responsiveness of your applications and servers, make sure that you have sufficient resources in place for typical traffic and high load situations:

- If you have your own server, make sure it has enough CPU, memory, and disk space to handle viewer requests, based on your evaluation.
- If you use an Amazon EC2 instance as your backend server, make sure that the instance type has the appropriate resources to fulfill incoming requests. For more information, see [Instance Types](#) in the Amazon EC2 User Guide.

In addition, consider the following tuning steps to help avoid timeouts:

- If the Time to First Byte value that is returned by the curl command seems high, take steps to improve the performance of your application. Improving application responsiveness will in turn help reduce timeout errors.
- Tune database queries to make sure that they can handle high request volumes without slow performance.
- Set up [keep-alive \(persistent\)](#) connections on your backend server. This option helps to avoid latencies that occur when connections must be re-established for subsequent requests or users.

- If you use ELB as your origin, learn how you can reduce latency by reviewing the suggestions in the following Knowledge Center article: [How to troubleshoot ELB high latency](#).

If needed, adjust the CloudFront timeout value

If you have evaluated and addressed slow application performance, origin server capacity, and other issues, but viewers are still experiencing HTTP 504 errors, then you should consider changing the time that is specified in your web distribution for origin response timeout. To learn more, see [Origin Response Timeout](#).

Load Testing CloudFront

Traditional load testing methods don't work well with CloudFront because CloudFront uses DNS to balance loads across geographically dispersed edge locations and within each edge location. When a client requests content from CloudFront, the client receives a DNS response that includes a set of IP addresses. If you test by sending requests to just one of the IP addresses that DNS returns, you're testing only a small subset of the resources in one CloudFront edge location, which doesn't accurately represent actual traffic patterns. Depending on the volume of data requested, testing in this way may overload and degrade the performance of that small subset of CloudFront servers.

CloudFront is designed to scale for viewers that have different client IP addresses and different DNS resolvers across multiple geographic regions. To perform load testing that accurately assesses CloudFront performance, we recommend that you do all of the following:

- Send client requests from multiple geographic regions.
- Configure your test so each client makes an independent DNS request; each client will then receive a different set of IP addresses from DNS.
- For each client that is making requests, spread your client requests across the set of IP addresses that are returned by DNS, which ensures that the load is distributed across multiple servers in a CloudFront edge location.

Request and Response Behavior

The following sections explain how CloudFront processes viewer requests and forwards the requests to your Amazon S3 or custom origin, and how CloudFront processes responses from your origin, including how CloudFront processes and caches 4xx and 5xx HTTP status codes.

Topics

- [Request and Response Behavior for Amazon S3 Origins \(p. 263\)](#)
- [Request and Response Behavior for Custom Origins \(p. 269\)](#)
- [Request and Response Behavior for Origin Groups \(p. 282\)](#)
- [Adding Custom Headers to Origin Requests \(p. 282\)](#)
- [How CloudFront Processes Partial Requests for an Object \(Range GETs\) \(p. 285\)](#)
- [How CloudFront Processes HTTP 3xx Status Codes from Your Origin \(p. 285\)](#)
- [How CloudFront Processes and Caches HTTP 4xx and 5xx Status Codes from Your Origin \(p. 286\)](#)

Request and Response Behavior for Amazon S3 Origins

Topics

- [How CloudFront Processes HTTP and HTTPS Requests \(p. 263\)](#)
- [How CloudFront Processes and Forwards Requests to Your Amazon S3 Origin Server \(p. 263\)](#)
- [How CloudFront Processes Responses from Your Amazon S3 Origin Server \(p. 268\)](#)

How CloudFront Processes HTTP and HTTPS Requests

For Amazon S3 origins, CloudFront accepts requests in both HTTP and HTTPS protocols for objects in a CloudFront distribution by default. CloudFront then forwards the requests to your Amazon S3 bucket using the same protocol in which the requests were made.

For custom origins, when you create your distribution, you can specify how CloudFront accesses your origin: HTTP only, or matching the protocol that is used by the viewer. For more information about how CloudFront handles HTTP and HTTPS requests for custom origins, see [Protocols \(p. 277\)](#).

For information about how to restrict your web distribution so that end users can only access objects using HTTPS, see [Using HTTPS with CloudFront \(p. 121\)](#). (This option doesn't apply to RTMP distributions, which use the RTMP protocol.)

Note

The charge for HTTPS requests is higher than the charge for HTTP requests. For more information about billing rates, go to the [CloudFront pricing plan](#).

How CloudFront Processes and Forwards Requests to Your Amazon S3 Origin Server

This topic contains information about how CloudFront processes viewer requests and forwards the requests to your Amazon S3 origin.

Topics

- [Caching Duration and Minimum TTL \(p. 264\)](#)
- [Client IP Addresses \(p. 264\)](#)
- [Conditional GETs \(p. 264\)](#)
- [Cookies \(p. 265\)](#)
- [Cross-Origin Resource Sharing \(CORS\) \(p. 265\)](#)
- [GET Requests That Include a Body \(p. 265\)](#)
- [HTTP Methods \(p. 265\)](#)
- [HTTP Request Headers That CloudFront Removes or Updates \(p. 266\)](#)
- [Maximum Length of a Request and Maximum Length of a URL \(p. 266\)](#)
- [OCSP Stapling \(p. 266\)](#)
- [Protocols \(p. 266\)](#)
- [Query Strings \(p. 267\)](#)
- [Origin Connection Timeout and Attempts \(p. 267\)](#)
- [Origin Response Timeout \(p. 267\)](#)
- [Simultaneous Requests for the Same Object \(Traffic Spikes\) \(p. 267\)](#)

Caching Duration and Minimum TTL

For web distributions, to control how long your objects stay in a CloudFront cache before CloudFront forwards another request to your origin, you can:

- Configure your origin to add a `Cache-Control` or an `Expires` header field to each object.
- Specify a value for Minimum TTL in CloudFront cache behaviors.
- Use the default value of 24 hours.

For more information, see [Managing How Long Content Stays in an Edge Cache \(Expiration\) \(p. 236\)](#).

Client IP Addresses

If a viewer sends a request to CloudFront and does not include an `X-Forwarded-For` request header, CloudFront gets the IP address of the viewer from the TCP connection, adds an `X-Forwarded-For` header that includes the IP address, and forwards the request to the origin. For example, if CloudFront gets the IP address `192.0.2.2` from the TCP connection, it forwards the following header to the origin:

```
X-Forwarded-For: 192.0.2.2
```

If a viewer sends a request to CloudFront and includes an `X-Forwarded-For` request header, CloudFront gets the IP address of the viewer from the TCP connection, appends it to the end of the `X-Forwarded-For` header, and forwards the request to the origin. For example, if the viewer request includes `X-Forwarded-For: 192.0.2.4,192.0.2.3` and CloudFront gets the IP address `192.0.2.2` from the TCP connection, it forwards the following header to the origin:

```
X-Forwarded-For: 192.0.2.4,192.0.2.3,192.0.2.2
```

Note

The `X-Forwarded-For` header contains IPv4 addresses (such as `192.0.2.44`) and IPv6 addresses (such as `2001:0db8:85a3:0000:0000:8a2e:0370:7334`).

Conditional GETs

When CloudFront receives a request for an object that has expired from an edge cache, it forwards the request to the Amazon S3 origin either to get the latest version of the object or to get confirmation from

Amazon S3 that the CloudFront edge cache already has the latest version. When Amazon S3 originally sent the object to CloudFront, it included an `ETag` value and a `LastModified` value in the response. In the new request that CloudFront forwards to Amazon S3, CloudFront adds one or both of the following:

- An `If-Match` or `If-None-Match` header that contains the `ETag` value for the expired version of the object.
- An `If-Modified-Since` header that contains the `LastModified` value for the expired version of the object.

Amazon S3 uses this information to determine whether the object has been updated and, therefore, whether to return the entire object to CloudFront or to return only an HTTP 304 status code (not modified).

Cookies

Amazon S3 doesn't process cookies. If you configure a cache behavior to forward cookies to an Amazon S3 origin, CloudFront forwards the cookies, but Amazon S3 ignores them. All future requests for the same object, regardless if you vary the cookie, are served from the existing object in the cache.

Cross-Origin Resource Sharing (CORS)

If you want CloudFront to respect Amazon S3 cross-origin resource sharing settings, configure CloudFront to forward selected headers to Amazon S3. For more information, see [Caching Content Based on Request Headers](#) (p. 246).

GET Requests That Include a Body

If a viewer `GET` request includes a body, CloudFront returns an HTTP status code 403 (Forbidden) to the viewer.

HTTP Methods

If you configure CloudFront to process all of the HTTP methods that it supports, CloudFront accepts the following requests from viewers and forwards them to your Amazon S3 origin:

- `DELETE`
- `GET`
- `HEAD`
- `OPTIONS`
- `PATCH`
- `POST`
- `PUT`

CloudFront always caches responses to `GET` and `HEAD` requests. You can also configure CloudFront to cache responses to `OPTIONS` requests. CloudFront does not cache responses to requests that use the other methods.

If you use an Amazon S3 bucket as the origin for your distribution and if you use CloudFront origin access identities, `POST` requests aren't supported in some Amazon S3 Regions and `PUT` requests in those Regions require an additional header. For more information, see [Using an OAI in Amazon S3 Regions that Support Only Signature Version 4 Authentication](#) (p. 215).

If you want to use multi-part uploads to add objects to an Amazon S3 bucket, you must add a CloudFront origin access identity to your distribution and grant the origin access identity the needed permissions. For more information, see [Restricting Access to Amazon S3 Content by Using an Origin Access Identity](#) (p. 209).

Important

If you configure CloudFront to accept and forward to Amazon S3 all of the HTTP methods that CloudFront supports, you must create a CloudFront origin access identity to restrict access to your Amazon S3 content and grant the origin access identity the required permissions. For example, if you configure CloudFront to accept and forward these methods because you want to use `PUT`, you must configure Amazon S3 bucket policies or ACLs to handle `DELETE` requests appropriately so viewers can't delete resources that you don't want them to. For more information, see [Restricting Access to Amazon S3 Content by Using an Origin Access Identity](#) (p. 209).

For information about the operations supported by Amazon S3, see the [Amazon S3 documentation](#).

HTTP Request Headers That CloudFront Removes or Updates

CloudFront removes or updates some headers before forwarding requests to your Amazon S3 origin. For most headers this behavior is the same as for custom origins. For a full list of HTTP request headers and how CloudFront processes them, see [HTTP Request Headers and CloudFront Behavior \(Custom and S3 Origins\)](#) (p. 273).

Maximum Length of a Request and Maximum Length of a URL

The maximum length of a request, including the path, the query string (if any), and headers, is 20,480 bytes.

CloudFront constructs a URL from the request. The maximum length of this URL is 8192 bytes.

If a request or a URL exceeds these maximums, CloudFront returns HTTP status code 413, Request Entity Too Large, to the viewer, and then terminates the TCP connection to the viewer.

OCSP Stapling

When a viewer submits an HTTPS request for an object, either CloudFront or the viewer must confirm with the certificate authority (CA) that the SSL certificate for the domain has not been revoked. OCSP stapling speeds up certificate validation by allowing CloudFront to validate the certificate and to cache the response from the CA, so the client doesn't need to validate the certificate directly with the CA.

The performance improvement of OCSP stapling is more pronounced when CloudFront receives a lot of HTTPS requests for objects in the same domain. Each server in a CloudFront edge location must submit a separate validation request. When CloudFront receives a lot of HTTPS requests for the same domain, every server in the edge location soon has a response from the CA that it can "staple" to a packet in the SSL handshake; when the viewer is satisfied that the certificate is valid, CloudFront can serve the requested object. If your distribution doesn't get much traffic in a CloudFront edge location, new requests are more likely to be directed to a server that hasn't validated the certificate with the CA yet. In that case, the viewer separately performs the validation step and the CloudFront server serves the object. That CloudFront server also submits a validation request to the CA, so the next time it receives a request that includes the same domain name, it has a validation response from the CA.

Protocols

CloudFront forwards HTTP or HTTPS requests to the origin server based on the protocol of the viewer request, either HTTP or HTTPS.

Important

If your Amazon S3 bucket is configured as a website endpoint, you cannot configure CloudFront to use HTTPS to communicate with your origin because Amazon S3 doesn't support HTTPS connections in that configuration.

Query Strings

For web distributions, you can configure whether CloudFront forwards query string parameters to your Amazon S3 origin. For RTMP distributions, CloudFront does not forward query string parameters. For more information, see [Caching Content Based on Query String Parameters \(p. 241\)](#).

Origin Connection Timeout and Attempts

Origin connection timeout is the number of seconds that CloudFront waits when trying to establish a connection to the origin.

Origin connection attempts is the number of times that CloudFront attempts to connect to the origin.

Together, these settings determine how long CloudFront tries to connect to the origin before failing over to the secondary origin (in the case of an origin group) or returning an error response to the viewer. By default, CloudFront waits as long as 30 seconds (3 attempts of 10 seconds each) before attempting to connect to the secondary origin or returning an error response. You can reduce this time by specifying a shorter connection timeout, fewer attempts, or both.

For more information, see [Controlling Origin Timeouts and Attempts \(p. 233\)](#).

Origin Response Timeout

The *origin response timeout*, also known as the *origin read timeout* or *origin request timeout*, applies to both of the following:

- The amount of time, in seconds, that CloudFront waits for a response after forwarding a request to the origin.
- The amount of time, in seconds, that CloudFront waits after receiving a packet of a response from the origin and before receiving the next packet.

CloudFront behavior depends on the HTTP method of the viewer request:

- **GET** and **HEAD** requests – If the origin doesn't respond within 30 seconds or stops responding for 30 seconds, CloudFront drops the connection. If the specified number of [origin connection attempts \(p. 42\)](#) is more than 1, CloudFront tries again to get a complete response. CloudFront tries up to 3 times, as determined by the value of the *origin connection attempts* setting. If the origin doesn't respond during the final attempt, CloudFront doesn't try again until it receives another request for content on the same origin.
- **DELETE**, **OPTIONS**, **PATCH**, **PUT**, and **POST** requests – If the origin doesn't respond within 30 seconds, CloudFront drops the connection and doesn't try again to contact the origin. The client can resubmit the request if necessary.

You can't change the response timeout for an Amazon S3 origin (an S3 bucket that is *not* configured with static website hosting).

Simultaneous Requests for the Same Object (Traffic Spikes)

When a CloudFront edge location receives a request for an object and either the object isn't currently in the cache or the object has expired, CloudFront immediately sends the request to your Amazon S3 origin. If there's a traffic spike—if additional requests for the same object arrive at the edge location

before Amazon S3 responds to the first request—CloudFront pauses briefly before forwarding additional requests for the object to your origin. Typically, the response to the first request will arrive at the CloudFront edge location before the response to subsequent requests. This brief pause helps to reduce unnecessary load on Amazon S3. If additional requests are not identical because, for example, you configured CloudFront to cache based on request headers or query strings, CloudFront forwards all of the unique requests to your origin.

When the response from the origin includes a `Cache-Control: no-cache` header, CloudFront typically forwards the next request for the same object to the origin to determine whether the object has been updated. However, when there's a traffic spike and CloudFront pauses after forwarding the first request to your origin, multiple viewer requests might arrive before CloudFront receives a response from the origin. When CloudFront receives a response that contains a `Cache-Control: no-cache` header, it sends the object in the response to the viewer that made the original request and to all of the viewers that requested the object during the pause. After the response arrives from the origin, CloudFront forwards the next viewer request for the same object to the origin. In CloudFront access logs, the first request is identified as a `Miss` in the `x-edge-result-type` column, and all subsequent requests that CloudFront received during the pause are identified as a `Hit`. For more information about access log file format, see [Web Distribution Standard Log File Format \(p. 445\)](#).

How CloudFront Processes Responses from Your Amazon S3 Origin Server

This topic contains information about how CloudFront processes responses from your Amazon S3 origin.

Topics

- [Canceled Requests \(p. 268\)](#)
- [HTTP Response Headers That CloudFront Removes or Updates \(p. 268\)](#)
- [Maximum File Size \(p. 269\)](#)
- [Redirects \(p. 269\)](#)

Canceled Requests

If an object is not in the edge cache, and if a viewer terminates a session (for example, closes a browser) after CloudFront gets the object from your origin but before it can deliver the requested object, CloudFront does not cache the object in the edge location.

HTTP Response Headers That CloudFront Removes or Updates

CloudFront removes or updates the following header fields before forwarding the response from your Amazon S3 origin to the viewer:

- `Set-Cookie` – If you configure CloudFront to forward cookies, it will forward the `Set-Cookie` header field to clients. For more information, see [Caching Content Based on Cookies \(p. 244\)](#).
- `Trailer`
- `Transfer-Encoding` – If your Amazon S3 origin returns this header field, CloudFront sets the value to `chunked` before returning the response to the viewer.
- `Upgrade`
- `Via` – CloudFront sets the value to the following in the response to the viewer:

Via: `http-version alphanumeric-string.cloudfront.net` (CloudFront)

For example, if the client makes a request over HTTP/1.1, the value is something like the following:

Via: `1.1 1026589cc7887e7a0dc7827b4example.cloudfront.net` (CloudFront)

Maximum File Size

The maximum size of a response body that CloudFront will return to the viewer is 20 GB. This includes chunked transfer responses that don't specify the `Content-Length` header value.

Redirects

You can configure an Amazon S3 bucket to redirect all requests to another host name; this can be another Amazon S3 bucket or an HTTP server. If you configure a bucket to redirect all requests and if the bucket is the origin for a CloudFront distribution, we recommend that you configure the bucket to redirect all requests to a CloudFront distribution using either the domain name for the distribution (for example, `d111111abcdef8.cloudfront.net`) or an alternate domain name (a CNAME) that is associated with a distribution (for example, `example.com`). Otherwise, viewer requests bypass CloudFront, and the objects are served directly from the new origin.

Note

If you redirect requests to an alternate domain name, you must also update the DNS service for your domain by adding a CNAME record. For more information, see [Using Custom URLs for Files by Adding Alternate Domain Names \(CNAMEs\)](#) (p. 71).

Here's what happens when you configure a bucket to redirect all requests:

1. A viewer (for example, a browser) requests an object from CloudFront.
2. CloudFront forwards the request to the Amazon S3 bucket that is the origin for your distribution.
3. Amazon S3 returns an HTTP status code 301 (Moved Permanently) as well as the new location.
4. CloudFront caches the redirect status code and the new location, and returns the values to the viewer. CloudFront does not follow the redirect to get the object from the new location.
5. The viewer sends another request for the object, but this time the viewer specifies the new location that it got from CloudFront:
 - If the Amazon S3 bucket is redirecting all requests to a CloudFront distribution, using either the domain name for the distribution or an alternate domain name, CloudFront requests the object from the Amazon S3 bucket or the HTTP server in the new location. When the new location returns the object, CloudFront returns it to the viewer and caches it in an edge location.
 - If the Amazon S3 bucket is redirecting requests to another location, the second request bypasses CloudFront. The Amazon S3 bucket or the HTTP server in the new location returns the object directly to the viewer, so the object is never cached in a CloudFront edge cache.

Request and Response Behavior for Custom Origins

Topics

- [How CloudFront Processes and Forwards Requests to Your Custom Origin Server](#) (p. 269)
- [How CloudFront Processes Responses from Your Custom Origin Server](#) (p. 279)

How CloudFront Processes and Forwards Requests to Your Custom Origin Server

This topic contains information about how CloudFront processes viewer requests and forwards the requests to your custom origin.

Topics

- [Authentication](#) (p. 270)
- [Caching Duration and Minimum TTL](#) (p. 270)
- [Client IP Addresses](#) (p. 271)
- [Client-Side SSL Authentication](#) (p. 271)
- [Compression](#) (p. 271)
- [Conditional Requests](#) (p. 271)
- [Cookies](#) (p. 272)
- [Cross-Origin Resource Sharing \(CORS\)](#) (p. 272)
- [Encryption](#) (p. 272)
- [GET Requests That Include a Body](#) (p. 272)
- [HTTP Methods](#) (p. 272)
- [HTTP Request Headers and CloudFront Behavior \(Custom and S3 Origins\)](#) (p. 273)
- [HTTP Version](#) (p. 276)
- [Maximum Length of a Request and Maximum Length of a URL](#) (p. 276)
- [OCSP Stapling](#) (p. 276)
- [Persistent Connections](#) (p. 277)
- [Protocols](#) (p. 277)
- [Query Strings](#) (p. 277)
- [Origin Connection Timeout and Attempts](#) (p. 277)
- [Origin Response Timeout](#) (p. 278)
- [Simultaneous Requests for the Same Object \(Traffic Spikes\)](#) (p. 278)
- [User-Agent Header](#) (p. 278)

Authentication

For `DELETE`, `GET`, `HEAD`, `PATCH`, `POST`, and `PUT` requests, if you configure CloudFront to forward the `Authorization` header to your origin, you can configure your origin server to request client authentication.

For `OPTIONS` requests, you can configure your origin server to request client authentication only if you use the following CloudFront settings:

- Configure CloudFront to forward the `Authorization` header to your origin.
- Configure CloudFront to *not* cache the response to `OPTIONS` requests.

You can configure CloudFront to forward requests to your origin using either HTTP or HTTPS; for more information, see [Using HTTPS with CloudFront](#) (p. 121).

Caching Duration and Minimum TTL

For web distributions, to control how long your objects stay in a CloudFront cache before CloudFront forwards another request to your origin, you can:

- Configure your origin to add a `Cache-Control` or an `Expires` header field to each object.
- Specify a value for Minimum TTL in CloudFront cache behaviors.
- Use the default value of 24 hours.

For more information, see [Managing How Long Content Stays in an Edge Cache \(Expiration\)](#) (p. 236).

Client IP Addresses

If a viewer sends a request to CloudFront and does not include an `X-Forwarded-For` request header, CloudFront gets the IP address of the viewer from the TCP connection, adds an `X-Forwarded-For` header that includes the IP address, and forwards the request to the origin. For example, if CloudFront gets the IP address `192.0.2.2` from the TCP connection, it forwards the following header to the origin:

```
X-Forwarded-For: 192.0.2.2
```

If a viewer sends a request to CloudFront and includes an `X-Forwarded-For` request header, CloudFront gets the IP address of the viewer from the TCP connection, appends it to the end of the `X-Forwarded-For` header, and forwards the request to the origin. For example, if the viewer request includes `X-Forwarded-For: 192.0.2.4, 192.0.2.3` and CloudFront gets the IP address `192.0.2.2` from the TCP connection, it forwards the following header to the origin:

```
X-Forwarded-For: 192.0.2.4, 192.0.2.3, 192.0.2.2
```

Some applications, such as load balancers (including Elastic Load Balancing), web application firewalls, reverse proxies, intrusion prevention systems, and API Gateway, append the IP address of the CloudFront edge server that forwarded the request onto the end of the `X-Forwarded-For` header. For example, if CloudFront includes `X-Forwarded-For: 192.0.2.2` in a request that it forwards to ELB and if the IP address of the CloudFront edge server is `192.0.2.199`, the request that your EC2 instance receives contains the following header:

```
X-Forwarded-For: 192.0.2.2, 192.0.2.199
```

Note

The `X-Forwarded-For` header contains IPv4 addresses (such as `192.0.2.44`) and IPv6 addresses (such as `2001:0db8:85a3:0000:0000:8a2e:0370:7334`).

Client-Side SSL Authentication

CloudFront does not support client authentication with client-side SSL certificates. If an origin requests a client-side certificate, CloudFront drops the request.

Compression

For more information, see [Serving compressed files](#) (p. 116).

Conditional Requests

When CloudFront receives a request for an object that has expired from an edge cache, it forwards the request to the origin either to get the latest version of the object or to get confirmation from the origin that the CloudFront edge cache already has the latest version. Typically, when the origin last sent the object to CloudFront, it included an `ETag` value, a `LastModified` value, or both values in the response. In the new request that CloudFront forwards to the origin, CloudFront adds one or both of the following:

- An `If-Match` or `If-None-Match` header that contains the `ETag` value for the expired version of the object.
- An `If-Modified-Since` header that contains the `LastModified` value for the expired version of the object.

The origin uses this information to determine whether the object has been updated and, therefore, whether to return the entire object to CloudFront or to return only an HTTP 304 status code (not modified).

Cookies

You can configure CloudFront to forward cookies to your origin. For more information, see [Caching Content Based on Cookies](#) (p. 244).

Cross-Origin Resource Sharing (CORS)

If you want CloudFront to respect cross-origin resource sharing settings, configure CloudFront to forward the `Origin` header to your origin. For more information, see [Caching Content Based on Request Headers](#) (p. 246).

Encryption

You can require viewers to use HTTPS to send requests to CloudFront and require CloudFront to forward requests to your custom origin by using the protocol that is used by the viewer. For more information, see the following distribution settings:

- [Viewer Protocol Policy](#) (p. 49)
- [Origin Protocol Policy](#) (p. 45)

CloudFront forwards HTTPS requests to the origin server using the SSLv3, TLSv1.0, TLSv1.1, and TLSv1.2 protocols. For custom origins, you can choose the SSL protocols that you want CloudFront to use when communicating with your origin:

- If you're using the CloudFront console, choose protocols by using the **Origin SSL Protocols** check boxes. For more information, see [Creating a Distribution](#) (p. 37).
- If you're using the CloudFront API, specify protocols by using the `OriginSslProtocols` element. For more information, see [OriginSslProtocols](#) and [DistributionConfig](#) in the *Amazon CloudFront API Reference*.

If the origin is an Amazon S3 bucket, CloudFront always uses TLSv1.2.

Important

Other versions of SSL and TLS are not supported.

For more information about using HTTPS with CloudFront, see [Using HTTPS with CloudFront](#) (p. 121). For lists of the ciphers that CloudFront supports for HTTPS communication between viewers and CloudFront, and between CloudFront and your origin, see [Supported protocols and ciphers between viewers and CloudFront](#) (p. 128).

GET Requests That Include a Body

If a viewer `GET` request includes a body, CloudFront returns an HTTP status code 403 (Forbidden) to the viewer.

HTTP Methods

If you configure CloudFront to process all of the HTTP methods that it supports, CloudFront accepts the following requests from viewers and forwards them to your custom origin:

- `DELETE`
- `GET`
- `HEAD`
- `OPTIONS`
- `PATCH`

- `POST`
- `PUT`

CloudFront always caches responses to `GET` and `HEAD` requests. You can also configure CloudFront to cache responses to `OPTIONS` requests. CloudFront does not cache responses to requests that use the other methods.

For information about configuring whether your custom origin processes these methods, see the documentation for your origin.

Important

If you configure CloudFront to accept and forward to your origin all of the HTTP methods that CloudFront supports, configure your origin server to handle all methods. For example, if you configure CloudFront to accept and forward these methods because you want to use `POST`, you must configure your origin server to handle `DELETE` requests appropriately so viewers can't delete resources that you don't want them to. For more information, see the documentation for your HTTP server.

HTTP Request Headers and CloudFront Behavior (Custom and S3 Origins)

The following table lists HTTP request headers that you can forward to both custom and Amazon S3 origins (with the exceptions that are noted). For each header, the table includes information about the following:

- CloudFront behavior if you don't configure CloudFront to forward the header to your origin, which causes CloudFront to cache your objects based on header values.
- Whether you can configure CloudFront to cache objects based on header values for that header.

You can configure CloudFront to cache objects based on values in the `Date` and `User-Agent` headers, but we don't recommend it. These headers have many possible values, and caching based on their values would cause CloudFront to forward significantly more requests to your origin.

For more information about caching based on header values, see [Caching Content Based on Request Headers](#) (p. 246).

Header	Behavior If You Don't Configure CloudFront to Cache Based on Header Values	Caching Based on Header Values Is Supported
Other-defined headers	CloudFront forwards the headers to your origin.	Yes
<code>Accept</code>	CloudFront removes the header.	Yes
<code>Accept-Charset</code>	CloudFront removes the header.	Yes
<code>Accept-Encoding</code>	<p>If the value contains <code>gzip</code> or <code>br</code>, CloudFront forwards a normalized <code>Accept-Encoding</code> header to your origin.</p> <p>For more information, see Cache compressed objects (uses the <code>Accept-Encoding</code> header) (p. 86) and Serving compressed files (p. 116).</p>	Yes

Header	Behavior If You Don't Configure CloudFront to Cache Based on Header Values	Caching Based on Header Values Is Supported
Accept-Language	CloudFront removes the header.	Yes
Authorization	<ul style="list-style-type: none"> GET and HEAD requests – CloudFront removes the Authorization header field before forwarding the request to your origin. OPTIONS requests – CloudFront removes the Authorization header field before forwarding the request to your origin if you configure CloudFront to cache responses to OPTIONS requests. <p>CloudFront forwards the Authorization header field to your origin if you do not configure CloudFront to cache responses to OPTIONS requests.</p> <ul style="list-style-type: none"> DELETE, PATCH, POST, and PUT requests – CloudFront does not remove the header field before forwarding the request to your origin. 	Yes
Cache-Control	CloudFront forwards the header to your origin.	No
CloudFront-Forwarded-Proto	<p>CloudFront does not add the header before forwarding the request to your origin.</p> <p>For more information, see Configuring Caching Based on the Protocol of the Request (p. 249).</p>	Yes
CloudFront-Is-Desktop-Viewer	<p>CloudFront does not add the header before forwarding the request to your origin.</p> <p>For more information, see Configuring Caching Based on the Device Type (p. 248).</p>	Yes
CloudFront-Is-Mobile-Viewer	<p>CloudFront does not add the header before forwarding the request to your origin.</p> <p>For more information, see Configuring Caching Based on the Device Type (p. 248).</p>	Yes
CloudFront-Is-Tablet-Viewer	<p>CloudFront does not add the header before forwarding the request to your origin.</p> <p>For more information, see Configuring Caching Based on the Device Type (p. 248).</p>	Yes
CloudFront-Viewer-Country	CloudFront does not add the header before forwarding the request to your origin.	Yes
Connection	CloudFront replaces this header with Connection: Keep-Alive before forwarding the request to your origin.	No
Content-Length	CloudFront forwards the header to your origin.	No

Header	Behavior If You Don't Configure CloudFront to Cache Based on Header Values	Caching Based on Header Values Is Supported
Content-MD5	CloudFront forwards the header to your origin.	Yes
Content-Type	CloudFront forwards the header to your origin.	Yes
Cookie	If you configure CloudFront to forward cookies, it will forward the Cookie header field to your origin. If you don't, CloudFront removes the Cookie header field. For more information, see Caching Content Based on Cookies (p. 244) .	No
Date	CloudFront forwards the header to your origin.	Yes, but not recommended
Expect	CloudFront removes the header.	Yes
From	CloudFront forwards the header to your origin.	Yes
Host	CloudFront sets the value to the domain name of the origin that is associated with the requested object. You can't cache based on the Host header for Amazon S3 or MediaStore origins.	Yes (custom) No (S3 and MediaStore)
If-Match	CloudFront forwards the header to your origin.	Yes
If-Modified-Since	CloudFront forwards the header to your origin.	Yes
If-None-Match	CloudFront forwards the header to your origin.	Yes
If-Range	CloudFront forwards the header to your origin.	Yes
If-Unmodified-Since	CloudFront forwards the header to your origin.	Yes
Max-Forwards	CloudFront forwards the header to your origin.	No
Origin	CloudFront forwards the header to your origin.	Yes
Pragma	CloudFront forwards the header to your origin.	No
Proxy-Authenticate	CloudFront removes the header.	No
Proxy-Authorization	CloudFront removes the header.	No
Proxy-Connection	CloudFront removes the header.	No
Range	CloudFront forwards the header to your origin. For more information, see How CloudFront Processes Partial Requests for an Object (Range GETs) (p. 285) .	Yes, by default
Referer	CloudFront removes the header.	Yes
Request-Range	CloudFront forwards the header to your origin.	No
TE	CloudFront removes the header.	No
Trailer	CloudFront removes the header.	No

Header	Behavior If You Don't Configure CloudFront to Cache Based on Header Values	Caching Based on Header Values Is Supported
Transfer-Encoding	CloudFront forwards the header to your origin.	No
Upgrade	CloudFront removes the header, unless you've established a WebSocket connection.	No (except for WebSocket connections)
User-Agent	CloudFront replaces the value of this header field with <code>Amazon CloudFront</code> . If you want CloudFront to cache your content based on the device the user is using, see Configuring Caching Based on the Device Type (p. 248).	Yes, but not recommended
Via	CloudFront forwards the header to your origin.	Yes
Warning	CloudFront forwards the header to your origin.	Yes
X-Amz-Cf-Id	CloudFront adds the header to the viewer request before forwarding the request to your origin. The header value contains an encrypted string that uniquely identifies the request.	No
X-Edge-*	CloudFront removes all <code>X-Edge-*</code> headers.	No
X-Forwarded-For	CloudFront forwards the header to your origin. For more information, see Client IP Addresses (p. 271).	Yes
X-Forwarded-Proto	CloudFront removes the header.	No
X-Real-IP	CloudFront removes the header.	No

HTTP Version

CloudFront forwards requests to your custom origin using HTTP/1.1.

Maximum Length of a Request and Maximum Length of a URL

The maximum length of a request, including the path, the query string (if any), and headers, is 20,480 bytes.

CloudFront constructs a URL from the request. The maximum length of this URL is 8192 bytes.

If a request or a URL exceeds these maximums, CloudFront returns HTTP status code 413, Request Entity Too Large, to the viewer, and then terminates the TCP connection to the viewer.

OCSP Stapling

When a viewer submits an HTTPS request for an object, either CloudFront or the viewer must confirm with the certificate authority (CA) that the SSL certificate for the domain has not been revoked. OCSP stapling speeds up certificate validation by allowing CloudFront to validate the certificate and to cache the response from the CA, so the client doesn't need to validate the certificate directly with the CA.

The performance improvement of OCSP stapling is more pronounced when CloudFront receives numerous HTTPS requests for objects in the same domain. Each server in a CloudFront edge location

must submit a separate validation request. When CloudFront receives a lot of HTTPS requests for the same domain, every server in the edge location soon has a response from the CA that it can "staple" to a packet in the SSL handshake; when the viewer is satisfied that the certificate is valid, CloudFront can serve the requested object. If your distribution doesn't get much traffic in a CloudFront edge location, new requests are more likely to be directed to a server that hasn't validated the certificate with the CA yet. In that case, the viewer separately performs the validation step and the CloudFront server serves the object. That CloudFront server also submits a validation request to the CA, so the next time it receives a request that includes the same domain name, it has a validation response from the CA.

Persistent Connections

When CloudFront gets a response from your origin, it tries to maintain the connection for several seconds in case another request arrives during that period. Maintaining a persistent connection saves the time required to re-establish the TCP connection and perform another TLS handshake for subsequent requests.

For more information, including how to configure the duration of persistent connections, see [Origin Keep-alive Timeout \(p. 46\)](#) in the section [Values That You Specify When You Create or Update a Distribution \(p. 38\)](#).

Protocols

CloudFront forwards HTTP or HTTPS requests to the origin server based on the following:

- The protocol of the request that the viewer sends to CloudFront, either HTTP or HTTPS.
- The value of the **Origin Protocol Policy** field in the CloudFront console or, if you're using the CloudFront API, the `OriginProtocolPolicy` element in the `DistributionConfig` complex type. In the CloudFront console, the options are **HTTP Only**, **HTTPS Only**, and **Match Viewer**.

If you specify **HTTP Only** or **HTTPS Only**, CloudFront forwards requests to the origin server using the specified protocol, regardless of the protocol in the viewer request.

If you specify **Match Viewer**, CloudFront forwards requests to the origin server using the protocol in the viewer request. Note that CloudFront caches the object only once even if viewers make requests using both HTTP and HTTPS protocols.

Important

If CloudFront forwards a request to the origin using the HTTPS protocol, and if the origin server returns an invalid certificate or a self-signed certificate, CloudFront drops the TCP connection.

For information about how to update a distribution using the CloudFront console, see [Updating a Distribution \(p. 63\)](#). For information about how to update a distribution using the CloudFront API, go to [UpdateDistribution](#) in the *Amazon CloudFront API Reference*.

Query Strings

You can configure whether CloudFront forwards query string parameters to your origin. For more information, see [Caching Content Based on Query String Parameters \(p. 241\)](#).

Origin Connection Timeout and Attempts

Origin connection timeout is the number of seconds that CloudFront waits when trying to establish a connection to the origin.

Origin connection attempts is the number of times that CloudFront attempts to connect to the origin.

Together, these settings determine how long CloudFront tries to connect to the origin before failing over to the secondary origin (in the case of an origin group) or returning an error response to the viewer. By default, CloudFront waits as long as 30 seconds (3 attempts of 10 seconds each) before attempting to connect to the secondary origin or returning an error response. You can reduce this time by specifying a shorter connection timeout, fewer attempts, or both.

For more information, see [Controlling Origin Timeouts and Attempts \(p. 233\)](#).

Origin Response Timeout

The *origin response timeout*, also known as the *origin read timeout* or *origin request timeout*, applies to both of the following:

- The amount of time, in seconds, that CloudFront waits for a response after forwarding a request to the origin.
- The amount of time, in seconds, that CloudFront waits after receiving a packet of a response from the origin and before receiving the next packet.

CloudFront behavior depends on the HTTP method of the viewer request:

- **GET** and **HEAD** requests – If the origin doesn't respond or stops responding within the duration of the response timeout, CloudFront drops the connection. If the specified number of [origin connection attempts \(p. 42\)](#) is more than 1, CloudFront tries again to get a complete response. CloudFront tries up to 3 times, as determined by the value of the *origin connection attempts* setting. If the origin doesn't respond during the final attempt, CloudFront doesn't try again until it receives another request for content on the same origin.
- **DELETE**, **OPTIONS**, **PATCH**, **PUT**, and **POST** requests – If the origin doesn't respond within 30 seconds, CloudFront drops the connection and doesn't try again to contact the origin. The client can resubmit the request if necessary.

For more information, including how to configure the origin response timeout, see [Origin Response Timeout \(p. 45\)](#).

Simultaneous Requests for the Same Object (Traffic Spikes)

When a CloudFront edge location receives a request for an object and either the object isn't currently in the cache or the object has expired, CloudFront immediately sends the request to your origin. If there's a traffic spike—if additional requests for the same object arrive at the edge location before your origin responds to the first request—CloudFront pauses briefly before forwarding additional requests for the object to your origin. Typically, the response to the first request will arrive at the CloudFront edge location before the response to subsequent requests. This brief pause helps to reduce unnecessary load on your origin server. If additional requests are not identical because, for example, you configured CloudFront to cache based on request headers or cookies, CloudFront forwards all of the unique requests to your origin.

User-Agent Header

If you want CloudFront to cache different versions of your objects based on the device that a user is using to view your content, we recommend that you configure CloudFront to forward one or more of the following headers to your custom origin:

- `CloudFront-Is-Desktop-Viewer`
- `CloudFront-Is-Mobile-Viewer`
- `CloudFront-Is-SmartTV-Viewer`
- `CloudFront-Is-Tablet-Viewer`

Based on the value of the `User-Agent` header, CloudFront sets the value of these headers to `true` or `false` before forwarding the request to your origin. If a device falls into more than one category, more than one value might be `true`. For example, for some tablet devices, CloudFront might set both `CloudFront-Is-Mobile-Viewer` and `CloudFront-Is-Tablet-Viewer` to `true`. For more information about configuring CloudFront to cache based on request headers, see [Caching Content Based on Request Headers](#) (p. 246).

You can configure CloudFront to cache objects based on values in the `User-Agent` header, but we don't recommend it. The `User-Agent` header has many possible values, and caching based on those values would cause CloudFront to forward significantly more requests to your origin.

If you do not configure CloudFront to cache objects based on values in the `User-Agent` header, CloudFront adds a `User-Agent` header with the following value before it forwards a request to your origin:

```
User-Agent = Amazon CloudFront
```

CloudFront adds this header regardless of whether the request from the viewer includes a `User-Agent` header. If the request from the viewer includes a `User-Agent` header, CloudFront removes it.

How CloudFront Processes Responses from Your Custom Origin Server

This topic contains information about how CloudFront processes responses from your custom origin.

Topics

- [100-Continue Responses](#) (p. 279)
- [Caching](#) (p. 279)
- [Canceled Requests](#) (p. 280)
- [Content Negotiation](#) (p. 280)
- [Cookies](#) (p. 280)
- [Dropped TCP Connections](#) (p. 280)
- [HTTP Response Headers that CloudFront Removes or Replaces](#) (p. 280)
- [Maximum File Size](#) (p. 281)
- [Origin Unavailable](#) (p. 281)
- [Redirects](#) (p. 281)
- [Transfer Encoding](#) (p. 282)

100-Continue Responses

Your origin cannot send more than one 100-Continue response to CloudFront. After the first 100-Continue response, CloudFront expects an HTTP 200 OK response. If your origin sends another 100-Continue response after the first one, CloudFront will return an error.

Caching

- Ensure that the origin server sets valid and accurate values for the `Date` and `Last-Modified` header fields.
- If requests from viewers include the `If-Match` or `If-None-Match` request header fields, set the `ETag` response header field. If you do not specify an `ETag` value, CloudFront ignores subsequent `If-Match` or `If-None-Match` headers.
- CloudFront normally respects a `Cache-Control: no-cache` header in the response from the origin. For an exception, see [Simultaneous Requests for the Same Object \(Traffic Spikes\)](#) (p. 278).

Canceled Requests

If an object is not in the edge cache, and if a viewer terminates a session (for example, closes a browser) after CloudFront gets the object from your origin but before it can deliver the requested object, CloudFront does not cache the object in the edge location.

Content Negotiation

If your origin returns `Vary: *` in the response, and if the value of **Minimum TTL** for the corresponding cache behavior is **0**, CloudFront caches the object but still forwards every subsequent request for the object to the origin to confirm that the cache contains the latest version of the object. CloudFront doesn't include any conditional headers, such as `If-None-Match` or `If-Modified-Since`. As a result, your origin returns the object to CloudFront in response to every request.

If your origin returns `Vary: *` in the response, and if the value of **Minimum TTL** for the corresponding cache behavior is any other value, CloudFront processes the `Vary` header as described in [HTTP Response Headers that CloudFront Removes or Replaces \(p. 280\)](#).

Cookies

If you enable cookies for a cache behavior, and if the origin returns cookies with an object, CloudFront caches both the object and the cookies. Note that this reduces cacheability for an object. For more information, see [Caching Content Based on Cookies \(p. 244\)](#).

Dropped TCP Connections

If the TCP connection between CloudFront and your origin drops while your origin is returning an object to CloudFront, CloudFront behavior depends on whether your origin included a `Content-Length` header in the response:

- **Content-Length header** – CloudFront returns the object to the viewer as it gets the object from your origin. However, if the value of the `Content-Length` header doesn't match the size of the object, CloudFront doesn't cache the object.
- **Transfer-Encoding: Chunked** – CloudFront returns the object to the viewer as it gets the object from your origin. However, if the chunked response is not complete, CloudFront does not cache the object.
- **No Content-Length header** – CloudFront returns the object to the viewer and caches it, but the object may not be complete. Without a `Content-Length` header, CloudFront cannot determine whether the TCP connection was dropped accidentally or on purpose.

We recommend that you configure your HTTP server to add a `Content-Length` header to prevent CloudFront from caching partial objects.

HTTP Response Headers that CloudFront Removes or Replaces

CloudFront removes or updates the following header fields before forwarding the response from your origin to the viewer:

- **Set-Cookie** – If you configure CloudFront to forward cookies, it will forward the `Set-Cookie` header field to clients. For more information, see [Caching Content Based on Cookies \(p. 244\)](#).
- **Trailer**
- **Transfer-Encoding** – If your origin returns this header field, CloudFront sets the value to `chunked` before returning the response to the viewer.
- **Upgrade**
- **Vary** – Note the following:

- If you configure CloudFront to forward any of the device-specific headers to your origin (CloudFront-Is-Desktop-Viewer, CloudFront-Is-Mobile-Viewer, CloudFront-Is-SmartTV-Viewer, CloudFront-Is-Tablet-Viewer) and you configure your origin to return Vary:User-Agent to CloudFront, CloudFront returns Vary:User-Agent to the viewer. For more information, see [Configuring Caching Based on the Device Type \(p. 248\)](#).
- If you configure your origin to include either Accept-Encoding or Cookie in the Vary header, CloudFront includes the values in the response to the viewer.
- If you configure CloudFront to forward a whitelist of headers to your origin, and if you configure your origin to return the header names to CloudFront in the Vary header (for example, Vary:Accept-Charset, Accept-Language), CloudFront returns the Vary header with those values to the viewer.
- For information about how CloudFront processes a value of * in the Vary header, see [Content Negotiation \(p. 280\)](#).
- If you configure your origin to include any other values in the Vary header, CloudFront removes the values before returning the response to the viewer.
- Via – CloudFront sets the value to the following in the response to the viewer:

Via: `http-version alphanumeric-string.cloudfront.net` (CloudFront)

For example, if the client makes a request over HTTP/1.1, the value is something like the following:

Via: `1.1 1026589cc7887e7a0dc7827b4example.cloudfront.net` (CloudFront)

Maximum File Size

The maximum size of a response body that CloudFront will return to the viewer is 20 GB. This includes chunked transfer responses that don't specify the Content-Length header value.

Origin Unavailable

If your origin server is unavailable and CloudFront gets a request for an object that is in the edge cache but that has expired (for example, because the period of time specified in the Cache-Control max-age directive has passed), CloudFront either serves the expired version of the object or serves a custom error page. For more information about CloudFront behavior when you've configured custom error pages, see [How CloudFront Processes Errors When You Have Configured Custom Error Pages \(p. 287\)](#).

In some cases, an object that is seldom requested is evicted and is no longer available in the edge cache. CloudFront can't serve an object that has been evicted.

Redirects

If you change the location of an object on the origin server, you can configure your web server to redirect requests to the new location. After you configure the redirect, the first time a viewer submits a request for the object, CloudFront sends the request to the origin, and the origin responds with a redirect (for example, 302 Moved Temporarily). CloudFront caches the redirect and returns it to the viewer. CloudFront does not follow the redirect.

You can configure your web server to redirect requests to one of the following locations:

- The new URL of the object on the origin server. When the viewer follows the redirect to the new URL, the viewer bypasses CloudFront and goes straight to the origin. As a result, we recommend that you not redirect requests to the new URL of the object on the origin.
- The new CloudFront URL for the object. When the viewer submits the request that contains the new CloudFront URL, CloudFront gets the object from the new location on your origin, caches it at the edge location, and returns the object to the viewer. Subsequent requests for the object will be served by the

edge location. This avoids the latency and load associated with viewers requesting the object from the origin. However, every new request for the object will incur charges for two requests to CloudFront.

Transfer Encoding

CloudFront supports only the `chunked` value of the `Transfer-Encoding` header. If your origin returns `Transfer-Encoding: chunked`, CloudFront returns the object to the client as the object is received at the edge location, and caches the object in chunked format for subsequent requests.

If the viewer makes a `Range GET` request and the origin returns `Transfer-Encoding: chunked`, CloudFront returns the entire object to the viewer instead of the requested range.

We recommend that you use chunked encoding if the content length of your response cannot be predetermined. For more information, see [Dropped TCP Connections \(p. 280\)](#).

Request and Response Behavior for Origin Groups

Requests to an origin group work the same as requests to an origin that is not set up as an origin group, except when there is an origin failover. As with any other origin, when CloudFront receives a request and the content is already cached in an edge location, the content is served to viewers from the cache. When there's a cache miss and the origin is an origin group, viewer requests are forwarded to the primary origin in the origin group.

The request and response behavior for the primary origin is the same as it is for an origin that isn't in an origin group. For more information, see [Request and Response Behavior for Amazon S3 Origins \(p. 263\)](#) and [Request and Response Behavior for Custom Origins \(p. 269\)](#).

The following describes the behavior for origin failover when the primary origin returns specific HTTP status codes:

- HTTP 2xx status code (success): CloudFront caches the file and returns it to the viewer.
- HTTP 3xx status code (redirection): CloudFront returns the status code to the viewer.
- HTTP 4xx or 5xx status code (client/server error): If the returned status code has been configured for failover, CloudFront sends the same request to the secondary origin in the origin group.
- HTTP 4xx or 5xx status code (client/server error): If the returned status code has not been configured for failover, CloudFront returns the error to the viewer.

CloudFront fails over to the secondary origin only when the HTTP method of the viewer request is `GET`, `HEAD`, or `OPTIONS`. CloudFront does not fail over when the viewer sends a different HTTP method (for example `POST`, `PUT`, and so on).

When CloudFront sends a request to a secondary origin, the response behavior is the same as for a CloudFront origin that's not in an origin group.

For more information about origin groups, see [Optimizing High Availability with CloudFront Origin Failover \(p. 231\)](#).

Adding Custom Headers to Origin Requests

You can configure CloudFront to add custom headers to the requests that it sends to your origin. These custom headers enable you to send and gather information from your origin that you don't get with typical viewer requests. These headers can even be customized for each origin. CloudFront supports custom headers for both for custom and Amazon S3 origins.

Note

Adding custom headers applies only to web distributions.

Topics

- [Use Cases for Origin Custom Headers \(p. 283\)](#)
- [Configuring CloudFront to Add Custom Headers to Origin Requests \(p. 283\)](#)
- [Custom Headers that CloudFront Can't Add to Origin Requests \(p. 284\)](#)
- [Configure CloudFront to Forward Authorization Headers \(p. 284\)](#)

Use Cases for Origin Custom Headers

You can use custom headers for a variety of things, such as the following:

Identifying requests from CloudFront

You can identify the requests that your origin receives from CloudFront. This can be useful if you want to know if users are bypassing CloudFront, or if you're using more than one CDN and you want information about which requests are coming from each CDN.

Note

If you're using an Amazon S3 origin and you enable [Amazon S3 server access logging](#), the logs don't include header information.

Determining which requests come from a particular distribution

If you configure more than one CloudFront distribution to use the same origin, you can add different custom headers in each distribution. You can then use the logs from your origin to determine which requests came from which CloudFront distribution.

Enabling cross-origin resource sharing (CORS)

If some of your viewers don't support cross-origin resource sharing (CORS), you can configure CloudFront to always add the `Origin` header to requests that it sends to your origin. Then you can configure your origin to return the `Access-Control-Allow-Origin` header for every request. You must also [configure CloudFront to respect CORS settings \(p. 248\)](#).

Controlling access to content

You can use custom headers to control access to content. By configuring your origin to respond to requests only when they include a custom header that gets added by CloudFront, you prevent users from bypassing CloudFront and accessing your content directly on the origin. For more information, see [Restricting Access to Files on Custom Origins \(p. 148\)](#).

Configuring CloudFront to Add Custom Headers to Origin Requests

To configure a web distribution to add custom headers to requests that it sends to your origin, update the origin configuration using one of the following methods:

- **CloudFront console** – When you create or update a distribution, specify header names and values in the **Origin Custom Headers** settings. For more information, see [Creating a Distribution \(p. 37\)](#) or [Updating a Distribution \(p. 63\)](#).
- **CloudFront API** – For each origin that you want to add custom headers to, specify the header names and values in the `CustomHeaders` field inside `Origin`. For more information, see [CreateDistribution](#) or [UpdateDistribution](#).

If the header names and values that you specify are not already present in the viewer request, CloudFront adds them to the origin request. If a header is present, CloudFront overwrites the header value before forwarding the request to the origin.

For the quotas (formerly known as limits) that apply to origin custom headers, see [Quotas on Custom Headers \(Web Distributions Only\)](#) (p. 499).

Custom Headers that CloudFront Can't Add to Origin Requests

You can't configure CloudFront to add any of the following headers to requests that it sends to your origin:

- Cache-Control
- Connection
- Content-Length
- Cookie
- Host
- If-Match
- If-Modified-Since
- If-None-Match
- If-Range
- If-Unmodified-Since
- Max-Forwards
- Pragma
- Proxy-Authorization
- Proxy-Connection
- Range
- Request-Range
- TE
- Trailer
- Transfer-Encoding
- Upgrade
- Via
- Headers that begin with X-Amz-
- Headers that begin with X-Edge-
- X-Real-Ip

Configure CloudFront to Forward Authorization Headers

When it forwards requests to your origin, CloudFront removes some viewer headers by default, including authorization headers. To always forward authorization headers for a specific cache behavior, whitelist the headers for that cache behavior. To learn more, see [Whitelist Headers](#) (p. 50).

How CloudFront Processes Partial Requests for an Object (Range GETs)

For a large object, an end user's browser or client might make multiple `GET` requests and use the `Range` request header to download the object in smaller units. These requests for ranges of bytes, sometimes known as `Range GET` requests, improve the efficiency of partial downloads and the recovery from partially failed transfers.

When CloudFront receives a `Range GET` request, it checks the cache in the edge location that received the request. If the cache in that edge location already contains the entire object or the requested portion of the object, CloudFront immediately serves the requested range from the cache.

If the cache doesn't contain the requested range, CloudFront forwards the request to the origin. (To optimize performance, CloudFront may request a larger range than the client requested in the `Range GET`.) What happens next depends on whether the origin supports `Range GET` requests:

- **If the origin supports `Range GET` requests:** It returns the requested range. CloudFront serves the requested range and also caches it for future requests. (Amazon S3 supports `Range GET` requests, as do some HTTP servers, for example, Apache and IIS. For information about whether your HTTP server does, see the documentation for your HTTP server.)
- **If the origin doesn't support `Range GET` requests:** It returns the entire object. CloudFront serves the current request by sending the entire object while also caching it for future requests. After CloudFront caches the entire object in an edge cache, it responds to new `Range GET` requests by serving the requested range.

In either case, CloudFront begins to serve the requested range or object to the end user as soon as the first byte arrives from the origin.

Note

If the viewer makes a `Range GET` request and the origin returns `Transfer-Encoding: chunked`, CloudFront returns the entire object to the viewer instead of the requested range.

CloudFront generally follows the RFC specification for the `Range` header. However, if your `Range` headers don't adhere to the following requirements, CloudFront returns HTTP status code 200 with the full object instead of status code 206 with the specified ranges:

- The ranges must be listed in ascending order. For example, `100–200, 300–400` is valid, `300–400, 100–200` is not valid.
- The ranges must not overlap. For example, `100–200, 150–250` is not valid.
- All of the ranges specifications must be valid. For example, you can't specify a negative value as part of a range.

For more information about the `Range` request header, see "Section 14.35 Range" in *Hypertext Transfer Protocol -- HTTP/1.1* at <http://www.w3.org/Protocols/rfc2616/rfc2616-sec14.html#sec14.35>.

How CloudFront Processes HTTP 3xx Status Codes from Your Origin

When CloudFront requests an object from your Amazon S3 bucket or custom origin server, your origin sometimes returns an HTTP 3xx status code. This typically indicates one of the following:

- The object's URL has changed (for example, status codes 301, 302, 307, or 308)
- The object hasn't changed since the last time CloudFront requested it (status code 304)

CloudFront caches 3xx responses according to the settings in your CloudFront distribution and the headers in the response. For more information, see [Managing How Long Content Stays in an Edge Cache \(Expiration\)](#) (p. 236).

If your origin returns a redirect status code (for example, 301 or 307), CloudFront doesn't follow the redirect. CloudFront passes along the 301 or 307 response to the viewer, who can follow the redirect by sending a new request.

How CloudFront Processes and Caches HTTP 4xx and 5xx Status Codes from Your Origin

Topics

- [How CloudFront Processes Errors When You Have Configured Custom Error Pages](#) (p. 287)
- [How CloudFront Processes Errors When You Have Not Configured Custom Error Pages](#) (p. 288)
- [HTTP 4xx and 5xx Status Codes that CloudFront Caches](#) (p. 289)

When CloudFront requests an object from your Amazon S3 bucket or custom origin server, your origin sometimes returns an HTTP 4xx or 5xx status code, which indicates that an error has occurred. CloudFront behavior depends on:

- Whether you have configured custom error pages.
- Whether you have configured how long you want CloudFront to cache error responses from your origin (error caching minimum TTL).
- The status code.
- For 5xx status codes, whether the requested object is currently in the CloudFront edge cache.
- For some 4xx status codes, whether the origin returns a `Cache-Control max-age` or `Cache-Control s-maxage` header.

CloudFront always caches responses to `GET` and `HEAD` requests. You can also configure CloudFront to cache responses to `OPTIONS` requests. CloudFront does not cache responses to requests that use the other methods.

If the origin doesn't respond, the CloudFront request to the origin times out which is considered an HTTP 5xx error from the origin, even though the origin didn't respond with that error. In that scenario, CloudFront continues to serve cached content. For more information, see [Origin Unavailable](#) (p. 281).

If you have enabled logging, CloudFront writes the results to the logs regardless of the HTTP status code.

For more information about features and options that relate to the error message returned from CloudFront, see the following:

- For information about settings for custom error pages in the CloudFront console, see [Custom Error Pages and Error Caching](#) (p. 60).
- For information about the error caching minimum TTL in the CloudFront console, see [Error Caching Minimum TTL \(seconds\)](#) (p. 60).
- For a list of the HTTP status codes that CloudFront caches, see [HTTP 4xx and 5xx Status Codes that CloudFront Caches](#) (p. 289).

How CloudFront Processes Errors When You Have Configured Custom Error Pages

If you have configured custom error pages, CloudFront behavior depends on whether the requested object is in the edge cache.

The Requested Object Is Not in the Edge Cache

CloudFront continues to try to get the requested object from your origin when all of the following are true:

- A viewer requests an object.
- The object isn't in the edge cache.
- Your origin returns an HTTP 4xx or 5xx status code and one of the following is true:
 - Your origin returns an HTTP 5xx status code instead of returning a 304 status code (Not Modified) or an updated version of the object.
 - Your origin returns an HTTP 4xx status code that is not restricted by a cache control header and is included in the following list of status codes: [HTTP 4xx and 5xx Status Codes that CloudFront Always Caches \(p. 289\)](#).
 - Your origin returns an HTTP 4xx status code without a `Cache-Control max-age` header or a `Cache-Control s-maxage` header, and the status code is included in the following list of status codes: [Control HTTP 4xx Status Codes that CloudFront Caches Based on Cache Control Headers \(p. 290\)](#).

CloudFront does the following:

1. In the CloudFront edge cache that received the viewer request, CloudFront checks your distribution configuration and gets the path of the custom error page that corresponds with the status code that your origin returned.
2. CloudFront finds the first cache behavior in your distribution that has a path pattern that matches the path of the custom error page.
3. The CloudFront edge location sends a request for the custom error page to the origin that is specified in the cache behavior.
4. The origin returns the custom error page to the edge location.
5. CloudFront returns the custom error page to the viewer that made the request, and also caches the custom error page for the maximum of the following:
 - The amount of time specified by the error caching minimum TTL (one minute by default)
 - The amount of time specified by a `Cache-Control max-age` header or a `Cache-Control s-maxage` header that is returned by the origin when the first request generated the error
6. After the caching time (determined in Step 5) has elapsed, CloudFront tries again to get the requested object by forwarding another request to your origin. CloudFront continues to retry at intervals specified by the error caching minimum TTL.

The Requested Object Is in the Edge Cache

CloudFront continues to serve the object that is currently in the edge cache when all of the following are true:

- A viewer requests an object.
- The object is in the edge cache but it has expired.

- Your origin returns an HTTP 5xx status code instead of returning a 304 status code (Not Modified) or an updated version of the object.

CloudFront does the following:

1. If your origin returns a 5xx status code, CloudFront serves the object even though it has expired. For the duration of the error caching minimum TTL, CloudFront continues to respond to viewer requests by serving the object from the edge cache.

If your origin returns a 4xx status code, CloudFront returns the status code, not the requested object, to the viewer.

2. After the error caching minimum TTL has elapsed, CloudFront tries again to get the requested object by forwarding another request to your origin. Note that if the object is not requested frequently, CloudFront might evict it from the edge cache while your origin server is still returning 5xx responses. For information about how long objects stay in CloudFront edge caches, see [Managing How Long Content Stays in an Edge Cache \(Expiration\)](#) (p. 236).

How CloudFront Processes Errors When You Have Not Configured Custom Error Pages

If you have not configured custom error pages, CloudFront behavior depends on whether the requested object is in the edge cache.

The Requested Object Is Not in the Edge Cache

CloudFront continues to try to get the requested object from your origin when all of the following are true:

- A viewer requests an object.
- The object isn't in the edge cache.
- Your origin returns an HTTP 4xx or 5xx status code and one of the following is true:
 - Your origin returns an HTTP 5xx status code instead of returning a 304 status code (Not Modified) or an updated version of the object.
 - Your origin returns an HTTP 4xx status code that is not restricted by a cache control header and is included in the following list of status codes: [HTTP 4xx and 5xx Status Codes that CloudFront Always Caches](#) (p. 289)
 - Your origin returns an HTTP 4xx status code without a `Cache-Control max-age` header or a `Cache-Control s-maxage` header and the status code is included in the following list of status codes: [Control HTTP 4xx Status Codes that CloudFront Caches Based on Cache Control Headers](#) (p. 290).

CloudFront does the following:

1. CloudFront returns the 4xx or 5xx status code to the viewer, and also caches status code in the edge cache that received the request for the maximum of the following:
 - The amount of time specified by the error caching minimum TTL (one minute by default)
 - The amount of time specified by a `Cache-Control max-age` header or a `Cache-Control s-maxage` header that is returned by the origin when the first request generated the error
2. For the duration of the caching time (determined in Step 1), CloudFront responds to subsequent viewer requests for the same object with the cached 4xx or 5xx status code.

3. After the caching time (determined in Step 1) has elapsed, CloudFront tries again to get the requested object by forwarding another request to your origin. CloudFront continues to retry at intervals specified by the error caching minimum TTL.

The Requested Object Is in the Edge Cache

CloudFront continues to serve the object that is currently in the edge cache when all of the following are true:

- A viewer requests an object.
- The object is in the edge cache but it has expired.
- Your origin returns an HTTP 5xx status code instead of returning a 304 status code (Not Modified) or an updated version of the object.

CloudFront does the following:

1. If your origin returns a 5xx error code, CloudFront serves the object even though it has expired. For the duration of the error caching minimum TTL (one minute by default), CloudFront continues to respond to viewer requests by serving the object from the edge cache.

If your origin returns a 4xx status code, CloudFront returns the status code, not the requested object, to the viewer.

2. After the error caching minimum TTL has elapsed, CloudFront tries again to get the requested object by forwarding another request to your origin. Note that if the object is not requested frequently, CloudFront might evict it from the edge cache while your origin server is still returning 5xx responses. For information about how long objects stay in CloudFront edge caches, see [Managing How Long Content Stays in an Edge Cache \(Expiration\)](#) (p. 236).

HTTP 4xx and 5xx Status Codes that CloudFront Caches

CloudFront caches HTTP 4xx and 5xx status codes returned by your origin, depending on the specific status code that is returned and whether your origin returns specific headers in the response.

HTTP 4xx and 5xx Status Codes that CloudFront Always Caches

CloudFront always caches the following HTTP 4xx and 5xx status codes returned by your origin. If you have configured a custom error page for an HTTP status code, CloudFront caches the custom error page.

404	Not Found
405	Method Not Allowed
414	Request-URI Too Large
500	Internal Server Error
501	Not Implemented
502	Bad Gateway
503	Service Unavailable

504	Gateway Time-out
-----	------------------

HTTP 4xx Status Codes that CloudFront Caches Based on Cache Control Headers

CloudFront only caches the following HTTP 4xx status codes returned by your origin if your origin returns a `Cache-Control max-age` or `Cache-Control s-maxage` header. If you have configured a custom error page for one of these HTTP status codes—and your origin returns one of the cache control headers—CloudFront caches the custom error page.

400	Bad Request
403	Forbidden
412	Precondition Failed
415	Unsupported Media Type

Generating Custom Error Responses

If the objects that you're serving through CloudFront are unavailable for some reason, your web server typically returns an HTTP status code to CloudFront. For example, if a viewer specifies an invalid URL, your web server returns a 404 status code to CloudFront, and CloudFront returns that status code to the viewer. The viewer displays a brief and sparsely formatted default message similar to this:

```
Not Found: The requested URL /myfilename.html was not found on this server.
```

But you can display a custom error message instead, if you like. You also have several options for managing how CloudFront responds when there's an error. To specify options for custom error messages, you update your CloudFront distribution to specify those values. For more information, see [Custom Error Pages and Error Caching \(p. 60\)](#) in the topic [Values That You Specify When You Create or Update a Distribution \(p. 38\)](#).

Topics


- [Creating a Custom Error Page for Specific HTTP Status Codes \(p. 291\)](#)
- [Storing Objects and Custom Error Pages in Different Locations \(p. 293\)](#)
- [Changing Response Codes Returned by CloudFront \(p. 293\)](#)
- [Controlling How Long CloudFront Caches Errors \(p. 294\)](#)
- [How CloudFront Responds When a Custom Error Page Is Unavailable \(p. 294\)](#)
- [Pricing for Custom Error Pages \(p. 295\)](#)
- [Configuring Error Response Behavior \(p. 295\)](#)

Creating a Custom Error Page for Specific HTTP Status Codes

If you'd rather display a custom error message instead of the default message—for example, a page that uses the same formatting as the rest of your website—you can have CloudFront return to the viewer an object (such as an HTML file) that contains your custom error message.

To specify the specific file that you want to return and the errors for which the file should be returned, you update your CloudFront distribution to specify those values. For more information, see [Custom Error Pages and Error Caching \(p. 60\)](#) in the topic [Values That You Specify When You Create or Update a Distribution \(p. 38\)](#).

For example, the following is a customized error message:



You can specify a different object for each supported HTTP status code, or you can use the same object for all of the supported status codes. You can also choose to specify objects for some status codes and not for others.

The objects that you're serving through CloudFront can be unavailable for a variety of reasons. These fall into two broad categories:

- **Client errors** indicate a problem with the request. For example, an object with the specified name isn't available, or the user doesn't have the permissions required to get an object in your Amazon S3 bucket. When a client error occurs, the origin returns an HTTP status code in the 400 range to CloudFront.
- **Server errors** indicate a problem with the origin server. For example, the HTTP server is busy or unavailable. When a server error occurs, either your origin server returns an HTTP status code in the 500 range to CloudFront, or CloudFront doesn't get a response from your origin server for a certain period of time and assumes a 504 status code (gateway timeout).

The HTTP status codes for which CloudFront can return a custom error page include the following:

- 400, 403, 404, 405, 414, 416
- 500, 501, 502, 503, 504

Note

You can create a custom error page for HTTP status code 416 (Requested Range Not Satisfiable), and you can change the HTTP status code that CloudFront returns to viewers when your origin returns a status code 416 to CloudFront. (For more information, see [Changing Response Codes Returned by CloudFront \(p. 293\)](#).) However, CloudFront doesn't cache status code 416 responses, so you can specify a value for **Error Caching Minimum TTL** for status code 416, but CloudFront doesn't use it.

For a detailed explanation of how CloudFront handles error responses from your origin, see [How CloudFront Processes and Caches HTTP 4xx and 5xx Status Codes from Your Origin \(p. 286\)](#).

Storing Objects and Custom Error Pages in Different Locations

If you want to store your objects and your custom error pages in different locations, your distribution must include a cache behavior for which the following is true:

- The value of **Path Pattern** matches the path to your custom error messages. For example, suppose you saved custom error pages for 4xx errors in an Amazon S3 bucket in a directory named `/4xx-errors`. Your distribution must include a cache behavior for which the path pattern routes requests for your custom error pages to that location, for example, `/4xx-errors/*`.
- The value of **Origin** specifies the value of **Origin ID** for the origin that contains your custom error pages.

For more information, see [Cache Behavior Settings \(p. 46\)](#) in the topic [Values That You Specify When You Create or Update a Distribution \(p. 38\)](#).

Changing Response Codes Returned by CloudFront

You can choose the HTTP status code CloudFront returns along with a custom error page for a given HTTP status code. For example, if your origin returns a 500 status code to CloudFront, you might want CloudFront to return a custom error page and a 200 status code (OK) to the viewer. There are a variety of reasons that you might want CloudFront to return a status code to the viewer that is different from the one that your origin returned to CloudFront:

- Some internet devices (some firewalls and corporate proxies, for example) intercept HTTP 4xx and 5xx and prevent the response from being returned to the viewer. If you substitute 200, the response typically won't be intercepted.
- If you don't care about distinguishing among different client errors or server errors, you can specify **400** or **500** as the value that CloudFront returns for all 4xx or 5xx status codes.
- You might want to return a 200 status code (OK) and static website so your customers don't know that your website is down.

If you enable CloudFront access logs and you configure CloudFront to change the HTTP status code in the response, the value of the `sc-status` column in access logs will contain the status code that you specify. However, the value of the `x-edge-result-type` column will not be affected; it will still contain the result type of the response from the origin. For example, suppose you configure CloudFront to return a status code of 200 to the viewer when the origin returns 404 (Not Found) to CloudFront. When the origin responds to a request with a 404 status code, the value in the `sc-status` column in the access log will be 200, but the value in the `x-edge-result-type` column will be `Error`.

You can configure CloudFront to return any of the following HTTP status codes along with a custom error page:

- 200
- 400, 403, 404, 405, 414, 416
- 500, 501, 502, 503, 504

Controlling How Long CloudFront Caches Errors

By default, when your origin returns an HTTP 4xx or 5xx status code, CloudFront caches these error responses for one minute. CloudFront then submits the next request for the object to your origin to see whether the problem that caused the error has been resolved and the requested object is now available.

Note

You can create a custom error page for HTTP status code 416 (Requested Range Not Satisfiable), and you can change the HTTP status code that CloudFront returns to viewers when your origin returns a status code 416 to CloudFront. (For more information, see [Changing Response Codes Returned by CloudFront \(p. 293\)](#).) However, CloudFront doesn't cache status code 416 responses, so although you can specify a value for **Error Caching Minimum TTL** for status code 416, CloudFront doesn't use it.

You can specify the error-caching duration—the **Error Caching Minimum TTL**—for each 4xx and 5xx status code that CloudFront caches. For a procedure, see [Configuring Error Response Behavior \(p. 295\)](#). When you specify a duration, note the following:

- If you specify a short error-caching duration, CloudFront forwards more requests to your origin than if you specify a longer duration. For 5xx errors, this may aggravate the problem that originally caused your origin to return an error.
- When your origin returns an error for an object, CloudFront responds to requests for the object either with the error response or with your custom error page until the error-caching duration elapses. If you specify a long error-caching duration, CloudFront might continue to respond to requests with an error response or your custom error page for a long time after the object becomes available again.

If you want to control how long CloudFront caches errors for individual objects, you can configure your origin server to add the applicable header to the error response for that object:

- **If the origin adds a `Cache-Control max-age` or `Cache-Control s-maxage` directive, or an `Expires` header:** CloudFront caches error responses for the greater of the value in the header or the value of **Error Caching Minimum TTL**.

Be aware that `Cache-Control max-age` and `Cache-Control s-maxage` values cannot be greater than the **Maximum TTL** value set for the cache behavior for which the error page is being fetched.

- **If the origin adds other `Cache-Control` directives or adds no headers:** CloudFront caches error responses for the value of **Error Caching Minimum TTL**.

If the expiration time for a 4xx or 5xx status code for an object is longer than you want to wait, you can invalidate the status code by using the URL of the requested object. If your origin is returning an error response for multiple objects, you need to invalidate each object separately. For more information about invalidating objects, see [Invalidating Files \(p. 110\)](#).

How CloudFront Responds When a Custom Error Page Is Unavailable

If you configure CloudFront to return a custom error page for an HTTP status code but the custom error page isn't available, CloudFront returns to the viewer the status code that CloudFront received from the origin that contains the custom error pages. For example, suppose your custom origin returns a 500 status code and you have configured CloudFront to get a custom error page for a 500 status code from an Amazon S3 bucket. However, someone accidentally deleted the custom error page from your bucket. CloudFront will return an HTTP 404 status code (not found) to the viewer that requested the object.

Pricing for Custom Error Pages

When CloudFront returns a custom error page to a viewer, you pay the standard CloudFront charges for the custom error page, not the charges for the requested object. For more information about CloudFront charges, see [Amazon CloudFront Pricing](#).

Configuring Error Response Behavior

You can use either the CloudFront API or console to configure CloudFront error responses. For information about using the CloudFront API to configure error responses, go to [PUT Distribution Config](#) in the *Amazon CloudFront API Reference*, and see the `CustomErrorResponse` element.

To configure CloudFront error responses using the console

1. Create the custom error pages that you want CloudFront to return to viewers when your origin returns HTTP 4xx or 5xx errors. Save the pages in a location that is accessible to CloudFront.

We recommend that you store custom error pages in an Amazon S3 bucket even if you're using a custom origin. If you store custom error pages on an HTTP server and the server starts to return 5xx errors, CloudFront can't get the files that you want to return to viewers because the origin server is unavailable.

2. Confirm that you have granted CloudFront at least `read` permission to your custom error page objects.

For more information about Amazon S3 permissions, see [Access Control](#) in the *Amazon Simple Storage Service Developer Guide*. For information on using the Amazon S3 console to update permissions, go to the *Amazon Simple Storage Service Console User Guide*.

3. (Optional) Configure your origin server to add `Cache-Control` directives or an `Expires` header along with the error response for specific objects, if applicable. For more information, see [Controlling How Long CloudFront Caches Errors](#) (p. 294).
4. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
5. In the list of distributions, select the distribution to update and choose **Distribution Settings**.
6. Choose the **Error Pages** tab. Then either choose **Create Custom Error Response**, or choose an existing error code and choose **Edit**.

CloudFront > Edit Custom Error Response

Custom Error Response Settings

HTTP Error Code: 404: Not Found

Error Caching Minimum TTL (seconds): 30

Customize Error Response: ☒ Yes ☐ No

Response Page Path: /4xx-errors/404-not-found

HTTP Response Code: 200: OK

Cancel Yes, Edit

7. Enter the applicable values. For more information, see [Custom Error Pages and Error Caching \(p. 60\)](#).
8. If you configured CloudFront to return custom error pages, add or update the applicable cache behaviors. For more information, see [Storing Objects and Custom Error Pages in Different Locations \(p. 293\)](#).
9. To save your changes, choose **Yes, Edit**.

Video on Demand and Live Streaming Video with CloudFront

You can use CloudFront to deliver video on demand (VOD) or live streaming video using any HTTP origin. One way you can set up video workflows in the cloud is by using CloudFront together with [AWS Media Services](#).

Topics

- [About Streaming Video: Video on Demand and Live Streaming \(p. 297\)](#)
- [Delivering Video on Demand \(VOD\) with CloudFront \(p. 298\)](#)
- [Delivering Live Streaming Video with CloudFront and AWS Media Services \(p. 300\)](#)
- [Working with RTMP Distributions \(p. 305\)](#)

About Streaming Video: Video on Demand and Live Streaming

You must use an encoder to package video content before CloudFront can distribute the content. The packaging process creates *segments* that contain your audio, video, and captions content. It also generates manifest files, which describe in a specific order what segments to play and when. Common package formats are MPEG DASH, Apple HLS, Microsoft Smooth Streaming, and CMAF.

Video on demand (VOD) streaming

For video on demand (VOD) streaming, your video content is stored on a server and viewers can watch it at any time. To make an asset that viewers can stream, use an encoder, such as [AWS Elemental MediaConvert](#), to format and package your media files.

After your video is packaged into the right formats, you can store it on a server or in an Amazon S3 bucket, and then deliver it with CloudFront as viewers request it.

Live video streaming

For live video streaming, your video content is streamed real time as live events happen, or is set up as a 24x7 live channel. To create live outputs for broadcast and streaming delivery, use an encoder such as AWS Elemental MediaLive, to compress the video and format it for viewing devices.

After your video is encoded, you can store it in AWS Elemental MediaStore or convert it into different delivery formats by using AWS Elemental MediaPackage. Use either of these origins to set up a CloudFront distribution to deliver the content. For specific steps and guidance for creating distributions that work together with these services, see [Serving Video Using AWS Elemental MediaStore as the Origin \(p. 300\)](#) and [Serving Live Video Formatted with AWS Elemental MediaPackage \(p. 301\)](#).

Wowza and Unified Streaming also provide tools that you can use for streaming video with CloudFront. For more information about using Wowza with CloudFront, see [Bring your Wowza Streaming Engine license to CloudFront live HTTP streaming](#) on the Wowza documentation website. For information about

using Unified Streaming with CloudFront for VOD streaming, see [Amazon CloudFront](#) on the Unified Streaming documentation website.

Delivering Video on Demand (VOD) with CloudFront

To deliver video on demand (VOD) streaming with CloudFront, use the following services:

- Amazon S3 to store the content in its original format and to store the transcoded video.
- An encoder (such as AWS Elemental MediaConvert) to transcode the video into streaming formats.
- CloudFront to deliver the transcoded video to viewers. For Microsoft Smooth Streaming, see [Configuring Video on Demand for Microsoft Smooth Streaming \(p. 298\)](#).

To create a VOD solution with CloudFront

1. Upload your content to an Amazon S3 bucket. To learn more about working with Amazon S3, see [the Amazon Simple Storage Service Developer Guide](#).
2. Transcode your content by using a MediaConvert job. The job converts your video into the formats required by the players that your viewers use. You can also use the job to create assets that vary in resolution and bitrate. These assets are used for adaptive bitrate (ABR) streaming, which adjusts the viewing quality depending on the viewer's available bandwidth. MediaConvert stores the transcoded video in an S3 bucket.
3. Deliver your converted content by using a CloudFront distribution. Viewers can watch the content on any device, at any time.

Tip

You can explore how to use an AWS CloudFormation template to deploy a VOD AWS solution together with all the associated components. To see the steps for using the template, see [Automated Deployment](#) in the *Video on Demand on AWS* guide.

Configuring Video on Demand for Microsoft Smooth Streaming

You have the following options for using CloudFront to distribute video on demand (VOD) content that you've transcoded into the Microsoft Smooth Streaming format:

- Specify a web server that runs Microsoft IIS and supports Smooth Streaming as the origin for your distribution.
- Enable Smooth Streaming in the cache behaviors of a CloudFront distribution. Because you can use multiple cache behaviors in a distribution, you can use one distribution for Smooth Streaming media files as well as other content.

Important

If you specify a web server running Microsoft IIS as your origin, do *not* enable Smooth Streaming in the cache behaviors of your CloudFront distribution. CloudFront can't use a Microsoft IIS server as an origin if you enable Smooth Streaming as a cache behavior.

If you enable Smooth Streaming in a cache behavior (that is, you do not have a server that is running Microsoft IIS), note the following:

- You can still distribute other content using the same cache behavior if the content matches the value of **Path Pattern** for that cache behavior.
- CloudFront can use either an Amazon S3 bucket or a custom origin for Smooth Streaming media files. CloudFront cannot use a Microsoft IIS Server as an origin if you enable Smooth Streaming for the cache behavior.
- You cannot invalidate media files in the Smooth Streaming format. If you want to update files before they expire, you must rename them. For more information, see [Adding, Removing, or Replacing Content That CloudFront Distributes](#) (p. 104).

For information about Smooth Streaming clients, see [Smooth Streaming Primer](#) on the Microsoft documentation website.

To use CloudFront to distribute Smooth Streaming files when a Microsoft IIS web server isn't the origin

1. Transcode your media files into Smooth Streaming fragmented MP4 format.
2. Do one of the following:
 - **If you're using the CloudFront console:** When you create or update a web distribution, enable Smooth Streaming in one or more of the distribution's cache behaviors.
 - **If you're using the CloudFront API:** Add the `SmoothStreaming` element to the `DistributionConfig` complex type for one or more of the distribution's cache behaviors.
3. Upload the Smooth Streaming files to your origin.
4. Create either a `clientaccesspolicy.xml` or a `crossdomainpolicy.xml` file, and add it to a location that is accessible at the root of your distribution, for example, `https://d111111abcdef8.cloudfront.net/clientaccesspolicy.xml`. The following is an example policy:

```
<?xml version="1.0" encoding="utf-8"?>
<access-policy>
  <cross-domain-access>
    <policy>
      <allow-from http-request-headers="*">
        <domain uri="*" />
      </allow-from>
      <grant-to>
        <resource path="/" include-subpaths="true" />
      </grant-to>
    </policy>
  </cross-domain-access>
</access-policy>
```

For more information, see [Making a Service Available Across Domain Boundaries](#) on the Microsoft Developer Network website.

5. For links in your application (for example, a media player), specify the URL for the media file in the following format:

```
https://d111111abcdef8.cloudfront.net/video/presentation.ism/Manifest
```

Delivering Live Streaming Video with CloudFront and AWS Media Services

To use AWS Media Services with CloudFront to deliver live content to a global audience, follow the guidance included in this section.

Use [AWS Elemental MediaLive](#) to encode live video streams in real time. To encode a large video stream, MediaLive compresses it into smaller versions (*encodes*) that can be distributed to your viewers.

After you compress a live video stream, you can use either of the following two main options to prepare and serve the content:

- **Convert your content into required formats, and then serve it:** If you require content in multiple formats, use [AWS Elemental MediaPackage](#) to package the content for different device types. When you package the content, you can also implement extra features and add digital rights management (DRM) to prevent unauthorized use of your content. For step-by-step instructions for using CloudFront to serve content that MediaPackage formatted, see [Serving Live Video Formatted with AWS Elemental MediaPackage](#) (p. 301).
- **Store and serve your content using scalable origin:** If MediaLive encoded content in the formats required by all of the devices that your viewers use, use a highly scalable origin like [AWS Elemental MediaStore](#) to serve the content. For step-by-step instructions for using CloudFront to serve content that is stored in a MediaStore container, see [Serving Video Using AWS Elemental MediaStore as the Origin](#) (p. 300).

After you've set up your origin by using one of these options, you can distribute live streaming video to viewers by using CloudFront.

Tip

You can learn about an AWS solution that automatically deploys services for building a highly available real-time viewing experience. To see the steps to automatically deploy this solution, see [Live Streaming Automated Deployment](#).

Topics

- [Serving Video Using AWS Elemental MediaStore as the Origin](#) (p. 300)
- [Serving Live Video Formatted with AWS Elemental MediaPackage](#) (p. 301)

Serving Video Using AWS Elemental MediaStore as the Origin

If you have video stored in an [AWS Elemental MediaStore](#) container, you can create a CloudFront distribution to serve the content.

To get started, you grant CloudFront access to your MediaStore container. Then you create a CloudFront distribution and configure it to work with MediaStore.

To serve content from an AWS Elemental MediaStore container

1. Follow the procedure at [Allowing Amazon CloudFront to Access Your MediaStore Container](#), and then return to these steps to create your distribution.
2. Create a distribution with the following settings:

Origin Domain Name

The data endpoint that is assigned to your MediaStore container. From the dropdown list, choose the MediaStore container for your live video. The format of a MediaStore origin is Container-OriginEndpointURL. For example, mymediastore.data.mediastore.us-east-1.amazonaws.com. For more information, see [Origin Domain Name \(p. 40\)](#).

Origin Path

The folder structure in the MediaStore container where your objects are stored. For more information, see [Origin Path \(p. 42\)](#).

Origin Custom Headers

Add header names and values if you want CloudFront to include custom headers when it forwards requests to your origin.

Viewer Protocol Policy

Choose **Redirect HTTP to HTTPS**. For more information, see [Viewer Protocol Policy \(p. 49\)](#).

Object Caching

If the encoder that you use can't set cache controls on all objects, choose **Customize**. If your encoder can set cache controls on all objects, choose **Origin Cache Headers**.

Minimum TTL, Maximum TTL, and Default TTL

Set as appropriate for your caching needs and segment durations.

Error Caching Minimum TTL

Set to 5 seconds or less, to help prevent serving stale content.

For the other settings, you can set specific values based on other technical requirements or the needs of your business. For a list of all the options for web distributions and information about setting them, see [Values That You Specify When You Create or Update a Distribution \(p. 38\)](#).

3. After CloudFront provisions your distribution, edit the cache behavior to set up cross-origin resource sharing (CORS) for your origin:
 1. Select the distribution, and then choose **Distribution Settings**.
 2. Choose **Behaviors**, select your origin, and then choose **Edit**.
 3. Under **Cache Based on Selected Request Headers**, choose **Whitelist**, and then, under **Whitelist Headers**, select **Origin**.

To learn more about CORS, see *Configuring CloudFront to Respect Cross-Origin Resource Sharing (CORS) Settings* in [Caching Content Based on Request Headers \(p. 246\)](#).

4. For links in your application (for example, a media player), specify the name of the media file in the same format that you use for other objects that you're distributing using CloudFront.

Serving Live Video Formatted with AWS Elemental MediaPackage

If you formatted a live stream by using AWS Elemental MediaPackage, you can create a CloudFront distribution and configure cache behaviors to serve the live stream. The following process assumes that you have already [created a channel](#) and [added endpoints](#) for your live video using MediaPackage.

Note

Instead of using the following process, you can choose to automatically create a CloudFront distribution when you save a channel in MediaPackage. For more information, see [Creating a Distribution from AWS Elemental MediaPackage](#) in the *AWS Elemental MediaPackage User Guide*.

To create a CloudFront distribution for MediaPackage manually, follow these steps:

Steps

- [Step 1: Create and configure a CloudFront distribution](#) (p. 302)
- [Step 2: Add the other endpoints as origins](#) (p. 303)
- [Step 3: Configure cache behaviors for all endpoints](#) (p. 303)
- [Step 4: Use CloudFront to serve the live stream channel](#) (p. 305)

Step 1: Create and configure a CloudFront distribution

Complete the following procedure to set up a CloudFront distribution for the live video channel that you created with MediaPackage.

To create a distribution for your live video channel

1. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
2. Choose **Create Distribution**.
3. On the **Select a delivery method for your content** page, in the **Web** section, choose **Get Started**.
4. Choose the settings for the distribution, including the following:

Origin Domain Name

The origin where your MediaPackage live video channel and endpoints are. Choose the text field, then from the dropdown list, choose the MediaPackage channel for your live video. You can map one channel to several origin endpoints.

If you created your channel using another AWS account, type the origin URL value into the field. The origin must be an HTTPS URL.

For more information, see [Origin Domain Name](#) (p. 40) in the [Values That You Specify When You Create or Update a Distribution](#) (p. 38).

Origin Path

The path to the MediaPackage endpoint from where the content is served. When you choose an origin domain name, CloudFront populates the origin path.

If you used a channel from another AWS account for **Origin Domain Name**, the **Origin Path** field is not filled in for you. You must get the correct origin path from the other account so that you can enter it manually.

For more information about how an origin path works, see [Origin Path](#) (p. 42) in [Values That You Specify When You Create or Update a Distribution](#) (p. 38).

For the other distribution settings, set specific values based on other technical requirements or the needs of your business. For a list of all the options for distributions and information about setting them, see [Values That You Specify When You Create or Update a Distribution](#) (p. 38).

When you finish choosing the other distribution settings, choose **Create Distribution**.

5. Choose the distribution that you just created, then choose the **Behaviors** tab.
6. Choose the default cache behavior, and specify the correct cache behavior settings for the channel that you chose for the origin. Later, you'll add one or more additional origins and edit cache behavior settings for them.
7. Go to the [CloudFront Distributions](#) page.
8. Wait until the value of the **Status** column for your distribution has changed from **In Progress** to **Deployed**, indicating that CloudFront has created your distribution.

Step 2: Add the other endpoints as origins

Repeat the steps here to add each of your MediaPackage channel endpoints to your distribution.

To add other endpoints as origins

1. On the CloudFront console, choose the distribution that you created for your channel.
2. Choose the **Origins and Origin Groups** tab, then choose **Create Origin**.
3. For **Origin Domain Name**, in the dropdown list, choose a MediaPackage endpoint for your channel. CloudFront automatically completes the **Origin Path** field.
4. For the other settings, set the values based on other technical requirements or the needs of your business. For more information, see [Origin Settings \(p. 40\)](#) in [Values That You Specify When You Create or Update a Distribution \(p. 38\)](#).
5. Choose **Create**.

Step 3: Configure cache behaviors for all endpoints

For each endpoint, you must configure cache behaviors to add path patterns that route requests correctly. The path patterns that you specify depend on the video format that you're serving. The following procedure includes the path pattern information to use for Apple HLS, CMAF, DASH, and Microsoft Smooth Streaming formats.

You typically set up two cache behaviors for each endpoint:

- The parent manifest, which is the index to your files.
- The segments, which are the files of the video content.

To create a cache behavior for an endpoint

1. On the CloudFront console, choose the distribution that you created for your channel.
2. Choose the **Behaviors** tab, then choose **Create Behavior**.
3. In the **Cache Behavior Settings** section, populate **Path Pattern** with the first pattern indicated in the following guidance for each endpoint type. For example, for a DASH endpoint, type `*.mpd` for **Path Pattern**.

Path Patterns

For an HLS endpoint, create the following two cache behaviors:

- For parent and child manifests, use `*.m3u8`.
- For the content segments, use `*.ts`.

For a CMAF endpoint, create the following two cache behaviors:

- For parent and child manifests, use `*.m3u8`.

- For the content segments, use `*.mp4`.

For a DASH endpoint, create the following two cache behaviors:

- For the parent manifest, use `*.mpd`.
- For the content segments, use `*.mp4`.

For a Microsoft Smooth Streaming endpoint, only a manifest is served, so you create only one cache behavior: `index.ism/*`.

4. For each cache behavior, specify values for the following settings:

Viewer Protocol Policy

Choose **Redirect HTTP to HTTPS**.

Cache Based on Selected Request Headers

Choose **None (improves caching)**.

For more information about improving caching, see [Increasing the Proportion of Requests that Are Served from CloudFront Edge Caches \(Cache Hit Ratio\)](#) (p. 228).

Object Caching

MediaPackage sets default `Cache-Control` headers that ensure correct playback behavior. If you want to use those values, choose **Use Origin Cache Headers**. However, you can increase cache times for video segments. For more information about customizing the time that objects stay in the CloudFront cache, see [Object Caching](#) (p. 50) in the [Values That You Specify When You Create or Update a Distribution](#) (p. 38).

Minimum TTL

Set to 5 seconds or less, to help prevent serving stale content.

Query String Forwarding and Caching

Choose **Forward all, cache based on whitelist**.

Query String Whitelist

Specify `m` as the query string parameter that you want CloudFront to use as the basis for caching. The MediaPackage response always includes the tag `?m=###` to capture the modified time of the endpoint. If content is already cached with a different value for this tag, CloudFront requests a new manifest instead of serving the cached version.

If you're using the time-shifted viewing functionality in MediaPackage, specify `start` and `end` as additional query string parameters on the cache behavior for manifest requests (`*.m3u8`, `*.mpd`, and `index.ism/*`). This way, content is served that's specific to the requested time period in the manifest request. For more information about time-shifted viewing and formatting content start and end request parameters, see [Time-shifted Viewing](#) in the *AWS Elemental MediaPackage User Guide*.

If you're using the manifest filtering feature in MediaPackage, specify `aws.manifestfilter` as an additional query string parameter on the cache behavior for manifest requests (`*.m3u8`, `*.mpd`, and `index.ism/*`). This configures your distribution to forward the `aws.manifestfilter` query string to your MediaPackage origin, which is required for the manifest filtering feature to work. For more information, see [Manifest Filtering](#) in the *AWS Elemental MediaPackage User Guide*.

5. Choose **Create**.
6. If your endpoint is not a Microsoft Smooth Streaming endpoint, choose **Create Behavior**, and then repeat these steps to create a second cache behavior.

Step 4: Use CloudFront to serve the live stream channel

After you create the distribution, add the origins, and create the cache behaviors, you can serve the live stream channel using CloudFront. CloudFront routes requests from viewers to the correct MediaPackage endpoints based on the settings that you configured for the cache behaviors.

For links in your application (for example, a media player), specify the URL for the media file in the standard format for CloudFront URLs. For more information, see [Customizing the URL Format for Files in CloudFront](#) (p. 106).

Working with RTMP Distributions

Note

Adobe designated Flash as end-of-life at the end of 2020. As a result, Amazon CloudFront will no longer support Adobe Flash Media Server and will be deprecating Real-Time Messaging Protocol (RTMP) distributions by December 31, 2020. For more information, [read the full announcement on the Amazon CloudFront discussion forum](#).

This section describes how you configure and manage RTMP distributions. RTMP distributions stream media files using Adobe Media Server and the Adobe Real-Time Messaging Protocol (RTMP). For information about how to create an RTMP distribution, see [Task List for Streaming Media Files Using RTMP](#) (p. 307).

Topics

- [RTMP Distributions](#) (p. 305)
- [How RTMP Distributions Work](#) (p. 306)
- [Task List for Streaming Media Files Using RTMP](#) (p. 307)
- [Creating an RTMP Distribution Using the CloudFront Console](#) (p. 308)
- [Values that You Specify When You Create or Update an RTMP Distribution](#) (p. 309)
- [Values that CloudFront Displays in the Console When You Create or Update an RTMP Distribution](#) (p. 313)
- [Configuring the Media Player](#) (p. 314)
- [Using an Amazon S3 Bucket as the Origin for an RTMP Distribution](#) (p. 315)
- [Creating Multiple RTMP Distributions for an Origin Server](#) (p. 316)
- [Restricting Access Using Crossdomain.xml](#) (p. 316)
- [Error Codes for RTMP Distributions](#) (p. 316)
- [Troubleshooting RTMP Distributions](#) (p. 316)

RTMP Distributions

Note

Adobe designated Flash as end-of-life at the end of 2020. As a result, Amazon CloudFront will no longer support Adobe Flash Media Server and will be deprecating Real-Time Messaging Protocol (RTMP) distributions by December 31, 2020. For more information, [read the full announcement on the Amazon CloudFront discussion forum](#).

RTMP distributions stream media files using Adobe Media Server and the Adobe Real-Time Messaging Protocol (RTMP). An RTMP distribution must use an Amazon S3 bucket as the origin.

For information about the values you specify when you create an RTMP distribution, see [Working with RTMP Distributions](#) (p. 305). For information about creating an RTMP distribution, see [Task List for Streaming Media Files Using RTMP](#) (p. 307).

How RTMP Distributions Work

Note

Adobe designated Flash as end-of-life at the end of 2020. As a result, Amazon CloudFront will no longer support Adobe Flash Media Server and will be deprecating Real-Time Messaging Protocol (RTMP) distributions by December 31, 2020. For more information, [read the full announcement on the Amazon CloudFront discussion forum](#).

To stream media files using CloudFront, you provide two types of files to your end users:

- Your media files
- A media player, for example, JW Player, Flowplayer, or Adobe Flash

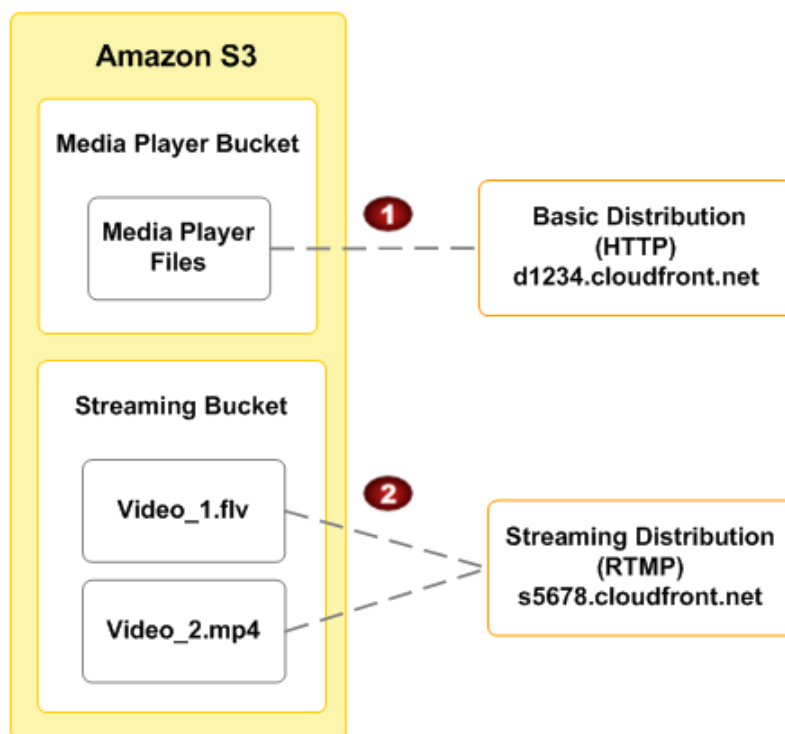
End users view your media files using the media player that you provide for them; they do not use the media player (if any) that is already installed on their computer or other device.

When an end user streams your media file, the media player begins to play the content of the file while the file is still being downloaded from CloudFront. The media file is not stored locally on the end user's system.

To use CloudFront to serve both the media player and the media files, you need two types of distributions: a web distribution for the media player, and an RTMP distribution for the media files. Web distributions serve files over HTTP, while RTMP distributions stream media files over RTMP (or a variant of RTMP).

The following example assumes that your media files and your media player are stored in different buckets in Amazon S3, but that isn't required—you can store media files and your media player in the same Amazon S3 bucket. You can also make the media player available to end users in other ways, for example, using CloudFront and a custom origin. However, the media files must use an Amazon S3 bucket as the origin.

In the following diagram, your site serves a cached copy of the media player to each end user through the `d1234.cloudfront.net` domain. The media player then accesses cached copies of your media files through the `s5678.cloudfront.net` domain.



1. Your media player bucket holds the media player and is the origin server for a regular HTTP distribution. In this example, the domain name for the distribution is `d1234.cloudfront.net`. (The `d` in `d1234.cloudfront.net` indicates that this is a web distribution.)
2. Your streaming media bucket holds your media files and is the origin server for an RTMP distribution. In this example, the domain name for the distribution is `s5678.cloudfront.net`. (The `s` in `s5678.cloudfront.net` indicates that this is an RTMP distribution.)

When you configure CloudFront to distribute media files, CloudFront uses Adobe Flash Media Server as the streaming server and streams your media files using Adobe's Real-Time Messaging Protocol (RTMP). CloudFront accepts RTMP requests over port 1935 and port 80.

CloudFront supports the following variants of the RTMP protocol:

- **RTMP** – Adobe's Real-Time Message Protocol
- **RTMPT** – Adobe streaming tunneled over HTTP
- **RTMPE** – Adobe encrypted
- **RTMPTE** – Adobe encrypted tunneled over HTTP

Task List for Streaming Media Files Using RTMP

Note

[Adobe designated Flash as end-of-life at the end of 2020](#). As a result, Amazon CloudFront will no longer support Adobe Flash Media Server and will be deprecating Real-Time Messaging Protocol (RTMP) distributions by December 31, 2020. For more information, [read the full announcement on the Amazon CloudFront discussion forum](#).

The following task list summarizes the process for creating a distribution for video on demand (VOD) streaming using the Adobe RTMP protocol for any media player.

To Create an RTMP Distribution

1. Create an Amazon S3 bucket for your media files. If you are using a different Amazon S3 bucket for your media player, create an Amazon S3 bucket for the media player files, too.

The names of your buckets must be all lowercase and cannot contain spaces.

2. Choose and configure a media player to play your media files. For more information, refer to the documentation for the media player.
3. Upload the files for your media player to the origin from which you want CloudFront to get the files. If you are using an Amazon S3 bucket as the origin for the media player, make the files (not the bucket) publicly readable.
4. Create a web distribution for your media player. (You can also use an existing distribution.) For more information, see [Steps for Creating a Distribution \(Overview\)](#) (p. 36).
5. Upload your media files to the Amazon S3 bucket that you created for the media files, and make the content (not the bucket) publicly readable.

Important

Media files in a Flash Video container must include the .flv filename extension, or the media will not stream.

You can put media player files and media files in the same bucket.

6. Create an RTMP distribution for your media files:
 - For more information about creating an RTMP distribution using the CloudFront console, see [Creating an RTMP Distribution Using the CloudFront Console](#) (p. 308).
 - For information about creating an RTMP distribution using the CloudFront API, see [CreateStreamingDistribution](#) in the *Amazon CloudFront API Reference*.
7. Configure your media player. For more information, see [Configuring the Media Player](#) (p. 314).

If you have trouble getting your content to play, see [Troubleshooting RTMP Distributions](#) (p. 316).

Creating an RTMP Distribution Using the CloudFront Console

Note

[Adobe designated Flash as end-of-life at the end of 2020](#). As a result, Amazon CloudFront will no longer support Adobe Flash Media Server and will be deprecating Real-Time Messaging Protocol (RTMP) distributions by December 31, 2020. For more information, [read the full announcement on the Amazon CloudFront discussion forum](#).

The following procedure explains how to create an RTMP distribution using the CloudFront console. If you want to create an RTMP distribution using the CloudFront API, go to [CreateStreamingDistribution](#) in the *Amazon CloudFront API Reference*.

For the current quota (formerly known as limit) on the number of RTMP distributions that you can create for each AWS account, see [Quotas](#) (p. 496). To request a higher quota, go to <https://console.aws.amazon.com/support/home#/case/create?issueType=service-limit-increase&limitType=service-code-cloudfront-distributions>.

To create an RTMP distribution using the CloudFront console

1. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
2. Click **Create Distribution**.
3. On the first page of the **Create Distribution Wizard**, in the **RTMP** section, choose **Get Started**.

4. Specify settings for the distribution. For more information, see [Values that You Specify When You Create or Update an RTMP Distribution \(p. 309\)](#).
5. Click **Create Distribution**.
6. After CloudFront creates your distribution, the value of the **Status** column for your distribution will change from **InProgress** to **Deployed**. If you chose to enable the distribution, it will then be ready to process requests. This should take less than 15 minutes.

The domain name that CloudFront assigns to your distribution appears in the list of distributions. The domain name also appears on the **General** tab for a selected distribution.

Values that You Specify When You Create or Update an RTMP Distribution

Note

Adobe designated Flash as end-of-life at the end of 2020. As a result, Amazon CloudFront will no longer support Adobe Flash Media Server and will be deprecating Real-Time Messaging Protocol (RTMP) distributions by December 31, 2020. For more information, [read the full announcement on the Amazon CloudFront discussion forum](#).

To stream media files using CloudFront, you create an RTMP distribution and specify the following values.

Topics

- [Origin Domain Name \(Amazon S3 Bucket\) \(p. 309\)](#)
- [Restrict Bucket Access \(Amazon S3 Only\) \(p. 310\)](#)
- [Origin Access Identity \(Amazon S3 Only\) \(p. 310\)](#)
- [Comment for New Identity\(Amazon S3 Only\) \(p. 310\)](#)
- [Your Identities \(Amazon S3 Only\) \(p. 310\)](#)
- [Grant Read Permissions on Bucket \(Amazon S3 Only\) \(p. 310\)](#)
- [Price Class \(p. 311\)](#)
- [Alternate Domain Names \(CNAMEs\) \(p. 311\)](#)
- [Logging \(p. 311\)](#)
- [Bucket for Logs \(p. 311\)](#)
- [Log Prefix \(p. 311\)](#)
- [Comment \(p. 312\)](#)
- [Distribution State \(p. 312\)](#)
- [Restrict Viewer Access \(Use Signed URLs\) \(p. 312\)](#)
- [Trusted Signers \(p. 312\)](#)
- [AWS Account Numbers \(p. 313\)](#)

Origin Domain Name (Amazon S3 Bucket)

The DNS domain name of the Amazon S3 bucket from which you want CloudFront to get objects for this origin, for example, `DOC-EXAMPLE-BUCKET1.s3.amazonaws.com`. In the CloudFront console, click in the **Origin Domain Name** field, and a list enumerates the Amazon S3 buckets that are associated with the current AWS account. To use a bucket from a different AWS account, type the domain name of the bucket in the following format:

`bucket-name.s3.region.amazonaws.com`

If you configured Amazon S3 Transfer Acceleration for your bucket, do not specify the `s3-accelerate` endpoint for **Origin Domain Name**.

The files must be publicly readable unless you secure your content in Amazon S3 by using a CloudFront origin access identity. For more information, see [Restricting Access to Amazon S3 Content by Using an Origin Access Identity](#) (p. 209).

Important

The bucket name must conform to DNS naming requirements. For more information, go to [Bucket Restrictions and Limitations](#) in the *Amazon Simple Storage Service Developer Guide*.

When you change the bucket from which CloudFront gets objects for the current origin, CloudFront immediately begins replicating the change to CloudFront edge locations. Until the distribution configuration is updated in a given edge location, CloudFront will continue to forward requests to the previous Amazon S3 bucket. As soon as the distribution configuration is updated in that edge location, CloudFront begins to forward requests to the new Amazon S3 bucket.

Changing the bucket does not require CloudFront to repopulate edge caches with objects from the new origin. As long as the viewer requests in your application have not changed, CloudFront will continue to serve objects that are already in an edge cache until the TTL on each object expires or until seldom-requested objects are evicted.

For more information, see [Using an Amazon S3 Bucket as the Origin for an RTMP Distribution](#) (p. 315).

Restrict Bucket Access (Amazon S3 Only)

Click **Yes** if you want to require end users to access objects in an Amazon S3 bucket by using only CloudFront URLs, not by using Amazon S3 URLs. Then specify the applicable values.

Click **No** if you want end users to be able to access objects using either CloudFront URLs or Amazon S3 URLs.

For more information, see [Restricting Access to Amazon S3 Content by Using an Origin Access Identity](#) (p. 209).

Origin Access Identity (Amazon S3 Only)

If you chose **Yes** for **Restrict Bucket Access**, choose whether to create a new origin access identity or use an existing one that is associated with your AWS account. If you already have an origin access identity, we recommend that you reuse it to simplify maintenance. For more information about origin access identities, see [Restricting Access to Amazon S3 Content by Using an Origin Access Identity](#) (p. 209).

Comment for New Identity(Amazon S3 Only)

If you chose **Create a New Identity** for **Origin Access Identity**, enter a comment that identifies the new origin access identity. CloudFront will create the origin access identity when you create this distribution.

Your Identities (Amazon S3 Only)

If you chose **Use an Existing Identity** for **Origin Access Identity**, choose the origin access identity that you want to use. You cannot use an origin access identity that is associated with another AWS account.

Grant Read Permissions on Bucket (Amazon S3 Only)

If you want CloudFront to automatically grant the origin access identity the permission to read objects in your Amazon S3 bucket, click **Yes, Update Bucket Policy**.

Important

If you click **Yes, Update Bucket Policy**, CloudFront updates the bucket policy to grant the specified origin access identity the permission to read objects in your bucket. However, CloudFront does not remove existing permissions in the bucket policy or permissions on individual objects. If users currently have permission to access the objects in your bucket using Amazon S3 URLs, they will still have that permission after CloudFront updates your bucket policy. To view or change the existing bucket policy and the existing permissions on the objects in your bucket, use a method provided by Amazon S3. For more information, see [Granting the OAI Permission to Read Files in Your Amazon S3 Bucket \(p. 212\)](#).

If you want to update permissions manually, for example, if you want to update ACLs on your objects instead of updating bucket permissions, click **No, I will Update Permissions**.

Price Class

The price class that corresponds with the maximum price that you want to pay for CloudFront service. By default, CloudFront serves your objects from edge locations in all CloudFront regions.

For more information about price classes and about how your choice of price class affects CloudFront performance for your distribution, see [Choosing the price class for a CloudFront distribution \(p. 10\)](#). For information about CloudFront pricing, including how price classes map to CloudFront regions, go to [Amazon CloudFront Pricing](#).

Alternate Domain Names (CNAMEs)

Optional. You can associate one or more CNAME aliases with a distribution so that you can use your domain name (for example, example.com) in the URLs for your objects instead of using the domain name that CloudFront assigned when you created your distribution. For more information, see [Using Custom URLs for Files by Adding Alternate Domain Names \(CNAMEs\) \(p. 71\)](#).

Logging

Whether you want CloudFront to log information about each request for an object and store the log files in an Amazon S3 bucket. You can enable or disable logging at any time. There is no extra charge if you enable logging, but you accrue the usual Amazon S3 charges for storing and accessing the files in an Amazon S3 bucket. You can delete the logs at any time. For more information about CloudFront access logs, see [Configuring and Using Standard Logs \(Access Logs\) \(p. 439\)](#).

Bucket for Logs

If you chose **On** for **Logging**, the Amazon S3 bucket that you want CloudFront to store access logs in, for example, `myLogs-DOC-EXAMPLE-BUCKET.s3.amazonaws.com`. If you enable logging, CloudFront records information about each end-user request for an object and stores the files in the specified Amazon S3 bucket. You can enable or disable logging at any time. For more information about CloudFront access logs, see [Configuring and Using Standard Logs \(Access Logs\) \(p. 439\)](#).

Note

You must have the permissions required to get and update Amazon S3 bucket ACLs, and the S3 ACL for the bucket must grant you `FULL_CONTROL`. This allows CloudFront to give the `awsdatafeeds` account permission to save log files in the bucket. For more information, see [Permissions Required to Configure Standard Logging and to Access Your Log Files \(p. 440\)](#).

Log Prefix

Optional. If you chose **On** for **Logging**, specify the string, if any, that you want CloudFront to prefix to the access log filenames for this distribution, for example, `exampleprefix/`. The trailing slash (/) is

optional but recommended to simplify browsing your log files. For more information about CloudFront access logs, see [Configuring and Using Standard Logs \(Access Logs\)](#) (p. 439).

Comment

Optional. When you create a distribution, you can include a comment of up to 128 characters. You can update the comment at any time.

Distribution State

When you create a distribution, you must specify whether you want the distribution to be enabled or disabled after it's created:

- *Enabled* means that as soon as the distribution is fully deployed you can deploy links that use the distribution's domain name and end users can retrieve content. Whenever a distribution is enabled, CloudFront accepts and processes any end-user requests for content that use the domain name associated with that distribution.

When you create, modify, or delete a CloudFront distribution, it takes time for your changes to propagate to the CloudFront database. An immediate request for information about a distribution might not show the change. Propagation usually completes within minutes, but a high system load or network partition might increase this time.

- *Disabled* means that even though the distribution might be deployed and ready to use, end users can't use it. When a distribution is disabled, CloudFront doesn't accept any end-user requests that use the domain name associated with that distribution. Until you switch the distribution from disabled to enabled (by updating the distribution's configuration), no one can use it.

You can toggle a distribution between disabled and enabled as often as you want. For information about updating a distribution's configuration, see [Updating a Distribution](#) (p. 63).

Restrict Viewer Access (Use Signed URLs)

If you want requests for objects served by this distribution to use public URLs, click **No**. If you want requests to use signed URLs, click **Yes**. Then specify the AWS accounts that you want to use to create signed URLs; these accounts are known as trusted signers.

For more information about trusted signers, see [Specifying the AWS Accounts That Can Create Signed URLs and Signed Cookies \(Trusted Signers\)](#) (p. 150).

Trusted Signers

Choose which AWS accounts you want to use as trusted signers for this distribution:

- **Self:** Use the account with which you're currently signed into the AWS Management Console as a trusted signer. If you're currently signed in as an IAM user, the associated AWS account is added as a trusted signer.
- **Specify Accounts:** Enter account numbers for trusted signers in the **AWS Account Numbers** field.

To create signed URLs, an AWS account must have at least one active CloudFront key pair.

Important

If you're updating a distribution that you're already using to distribute content, add trusted signers only when you're ready to start generating signed URLs for your objects. After you add trusted signers to a distribution, users must use signed URLs to access the objects served by this distribution.

AWS Account Numbers

If you want to create signed URLs using AWS accounts in addition to or instead of the current account, enter one AWS account number per line in this field. Note the following:

- The accounts that you specify must have at least one active CloudFront key pair. For more information, see [Creating CloudFront Key Pairs for Your Trusted Signers \(p. 151\)](#).
- You can't create CloudFront key pairs for IAM users, so you can't use IAM users as trusted signers.
- For information about how to get the AWS account number for an account, see [How Do I Get Security Credentials?](#) in the *Amazon Web Services General Reference*.
- If you enter the account number for the current account, CloudFront automatically checks the **Self** checkbox and removes the account number from the **AWS Account Numbers** list.

Values that CloudFront Displays in the Console When You Create or Update an RTMP Distribution

Note

Adobe designated Flash as end-of-life at the end of 2020. As a result, Amazon CloudFront will no longer support Adobe Flash Media Server and will be deprecating Real-Time Messaging Protocol (RTMP) distributions by December 31, 2020. For more information, [read the full announcement on the Amazon CloudFront discussion forum](#).

When you create a new RTMP distribution or update an existing distribution, CloudFront displays the following information in the CloudFront console.

Note

Active trusted signers, the AWS accounts that have an active CloudFront key pair and can be used to create valid signed URLs, are currently not visible in the CloudFront console.

Distribution ID

When you perform an action on a distribution using the CloudFront API, you use the distribution ID to specify which distribution you want to perform the action on, for example, `EDFDVBD6EXAMPLE`. You can't change the distribution ID.

Status

The possible status values for a distribution are listed in the following table.

Value	Description
InProgress	The distribution is still being created or updated.
Deployed	The distribution has been created or updated and the changes have been fully propagated through the CloudFront system.

In addition to ensuring that the status for a distribution is **Deployed**, you must enable the distribution before end users can use CloudFront to access your content. For more information, see [Distribution State \(p. 312\)](#).

Last Modified

The date and time that the distribution was last modified, using ISO 8601 format, for example, 2012-05-19T19:37:58Z. For more information, go to <http://www.w3.org/TR/NOTE-datetime>.

Domain Name

You use the distribution's domain name in the links to your objects, unless you're using alternate domain names (CNAMEs). For example, if your distribution's domain name is `d111111abcdef8.cloudfront.net`, the link to the example `/images/image.jpg` file would be `http://d111111abcdef8.cloudfront.net/images/image.jpg`. You can't change the CloudFront domain name for your distribution. For more information about CloudFront URLs for links to your objects, see [Customizing the URL Format for Files in CloudFront](#) (p. 106).

If you specified one or more alternate domain names (CNAMEs), you can use your own domain names for links to your objects instead of using the CloudFront domain name. For more information about CNAMEs, see [Alternate Domain Names \(CNAMEs\)](#) (p. 54).

Note

CloudFront domain names are unique. Your distribution's domain name was never used for a previous distribution and will never be reused for another distribution in the future.

Configuring the Media Player

Note

[Adobe designated Flash as end-of-life at the end of 2020](#). As a result, Amazon CloudFront will no longer support Adobe Flash Media Server and will be deprecating Real-Time Messaging Protocol (RTMP) distributions by December 31, 2020. For more information, [read the full announcement on the Amazon CloudFront discussion forum](#).

To play a media file, you must configure the media player with the correct path to the file. How you configure the media depends on which media player you're using and how you're using it.

When you configure the media player, the path you specify to the media file must contain the characters `cfx/st` immediately after the domain name, for example:

```
rtmp://s5c39gqb8ow64r.cloudfront.net/cfx/st/mediafile.flv.
```

Note

CloudFront follows Adobe's FMS naming requirements. Different players have their own rules about how to specify streams. The example above is for JW Player. Check your player's documentation. For example, Adobe's Flash Media Server does not allow the `.flv` extension to be present on the play path. Many players remove the `.flv` extension for you.

Your media player might ask for the path separate from the file name. For example, with JW Player, you specify a `streamer` and `file` variable:

- **streamer** – `rtmp://s5c39gqb8ow64r.cloudfront.net/cfx/st` (with no trailing slash)
- **file** – `mediafile.flv`

If you've stored the media files in a directory in your bucket (for example, `videos/mediafile.flv`), then the variables for JW Player would be:

- **streamer** – `rtmp://s5c39gqb8ow64r.cloudfront.net/cfx/st` (with no trailing slash)
- **file** – `videos/mediafile.flv`

For more information about JW Player, go to the [JW Player website](#).

MPEG Files

To serve MP3 audio files or H.264/MPEG-4 video files, you might need to prefix the file name with `mp3:` or `mp4:`. Some media players can be configured to add the prefix automatically. The media player might

also require you to specify the file name without the file extension (for example, `magicvideo` instead of `magicvideo.mp4`).

Using an Amazon S3 Bucket as the Origin for an RTMP Distribution

Note

[Adobe designated Flash as end-of-life at the end of 2020](#). As a result, Amazon CloudFront will no longer support Adobe Flash Media Server and will be deprecating Real-Time Messaging Protocol (RTMP) distributions by December 31, 2020. For more information, [read the full announcement on the Amazon CloudFront discussion forum](#).

When you create a distribution, you specify where CloudFront gets the files that it distributes to edge locations. For an RTMP distribution, you must use an Amazon S3 bucket; custom origins are not supported. To get your objects into your bucket, you can use any method supported by Amazon S3, for example, the Amazon S3 API or a third-party tool. You can create a hierarchy in your bucket just as you would with any other Amazon S3 bucket. You incur regular Amazon S3 charges for storing the objects in the bucket. For more information about the charges to use CloudFront, see [CloudFront Reports in the Console](#) (p. 405).

Using an existing Amazon S3 bucket as your CloudFront origin server doesn't change the bucket in any way; you can still use it as you normally would to store and access Amazon S3 objects (at the normal Amazon S3 prices).

You can use the same Amazon S3 bucket for both RTMP and web distributions.

Note

After you create an RTMP distribution, you can't change its origin server. If you need to change the Amazon S3 bucket for an RTMP distribution, you must create a new distribution that uses the new bucket and update either your links or your DNS records to use the domain name for the new distribution. You can then delete the original distribution. For more information, see [Deleting a Distribution](#) (p. 65).

When you specify the Amazon S3 bucket that you want CloudFront to get objects from, we recommend that you use the following format to access the bucket:

`bucket-name.s3.region.amazonaws.com`

Another option might be to use the following more general format, but be aware that this format doesn't work for Regions launched in 2019 or later:

`bucket-name.s3.amazonaws.com`

Do not specify the name of the bucket using the following values:

- The Amazon S3 path style, `s3.amazonaws.com/bucket-name`
- The Amazon S3 CNAME, if any

Important

For your bucket to work with CloudFront, the name must conform to DNS naming requirements. For more information, go to [Bucket Restrictions and Limitations](#) in the *Amazon Simple Storage Service Developer Guide*.

Creating Multiple RTMP Distributions for an Origin Server

Note

[Adobe designated Flash as end-of-life at the end of 2020](#). As a result, Amazon CloudFront will no longer support Adobe Flash Media Server and will be deprecating Real-Time Messaging Protocol (RTMP) distributions by December 31, 2020. For more information, [read the full announcement on the Amazon CloudFront discussion forum](#).

You typically create one RTMP distribution per Amazon S3 bucket, but you can choose to create multiple RTMP distributions for the same bucket. For example, if you had two distributions for an Amazon S3 bucket, you could reference a single media file using either distribution. In this case, if you had a media file called `media.flv` in your origin server, CloudFront would work with each distribution as though it referenced an individual `media.flv` object: one `media.flv` accessible through one distribution, and another `media.flv` accessible through the other distribution.

Restricting Access Using Crossdomain.xml

Note

[Adobe designated Flash as end-of-life at the end of 2020](#). As a result, Amazon CloudFront will no longer support Adobe Flash Media Server and will be deprecating Real-Time Messaging Protocol (RTMP) distributions by December 31, 2020. For more information, [read the full announcement on the Amazon CloudFront discussion forum](#).

The Adobe Flash Media Server `crossdomain.xml` file specifies which domains can access media files in a particular domain. CloudFront supplies a default file that allows all domains to access the media files in your RTMP distribution, and you cannot change this behavior. If you include a more restrictive `crossdomain.xml` file in your Amazon S3 bucket, CloudFront ignores it.

Error Codes for RTMP Distributions

Note

[Adobe designated Flash as end-of-life at the end of 2020](#). As a result, Amazon CloudFront will no longer support Adobe Flash Media Server and will be deprecating Real-Time Messaging Protocol (RTMP) distributions by December 31, 2020. For more information, [read the full announcement on the Amazon CloudFront discussion forum](#).

The following table lists the error codes that CloudFront can send to your media player. The errors are part of the string returned with `Event.info.application.message` or `Event.info.description`.

Error	Description
<code>DistributionNotFound</code>	The distribution was not found.
<code>DistributionTypeMismatch</code>	The distribution is not an RTMP distribution.
<code>InvalidInstance</code>	The instance is invalid.
<code>InvalidURI</code>	The URI is invalid.

Troubleshooting RTMP Distributions

Note

[Adobe designated Flash as end-of-life at the end of 2020](#). As a result, Amazon CloudFront will no longer support Adobe Flash Media Server and will be deprecating Real-Time Messaging

Protocol (RTMP) distributions by December 31, 2020. For more information, [read the full announcement on the Amazon CloudFront discussion forum](#).

If you're having trouble getting your media files to play, check the following items.

Item to Check	Description
Separate distributions for the media player files and media files	The media player must be served by a regular HTTP distribution (for example, domain name <code>d111111abcdef8.cloudfront.net</code>), and media files must be served by an RTMP distribution (for example, domain name <code>s5c39gqb8ow64r.cloudfront.net</code>). Make sure you're not using the same distribution for both.
<code>/cfx/st</code> in the file path	Confirm that the path for the file includes <code>/cfx/st</code> . You don't need to include <code>/cfx/st</code> in the path to the object in the Amazon S3 bucket. For more information, see Configuring the Media Player (p. 314) .
File names in the file path	<p>Some media players require that you include the file name extension (for example, <code>mp4:</code>) before the file name in the file path. Some media players also require that you exclude the file name extension (for example, <code>.mp4</code>) from the file path. For more information, see MPEG Files (p. 314).</p> <p>Note The names of the media files in your Amazon S3 bucket must always include the applicable file name extension.</p>
Port 1935 on your firewall	Adobe Flash Media Server uses port 1935 for RTMP. Make sure your firewall has this port open. If it doesn't, the typical message returned is "Unable to play video." You can also switch to RTMPT to tunnel over HTTP using port 80.
Adobe Flash Player messaging	<p>By default, the Adobe Flash Player won't display a message if the video file it's trying to play is missing. Instead, it waits for the file to show up. You might want to change this behavior to give your end users a better experience.</p> <p>To instead have the player send a message if the video is missing, use <code>play("vid", 0, -1)</code> instead of <code>play("vid")</code>.</p>

Customizing Content at the Edge with Lambda@Edge

Lambda@Edge is an extension of AWS Lambda, a compute service that lets you execute functions that customize the content that CloudFront delivers. You can author Node.js or Python functions in one Region, US-East-1 (N. Virginia), and then execute them in AWS locations globally that are closer to the viewer, without provisioning or managing servers. Lambda@Edge scales automatically, from a few requests per day to thousands per second. Processing requests at AWS locations closer to the viewer instead of on origin servers significantly reduces latency and improves the user experience.

When you associate a CloudFront distribution with a Lambda@Edge function, CloudFront intercepts requests and responses at CloudFront edge locations. You can execute Lambda functions when the following CloudFront events occur:

- When CloudFront receives a request from a viewer (viewer request)
- Before CloudFront forwards a request to the origin (origin request)
- When CloudFront receives a response from the origin (origin response)
- Before CloudFront returns the response to the viewer (viewer response)

There are many uses for Lambda@Edge processing. For example:

- A Lambda function can inspect cookies and rewrite URLs so that users see different versions of a site for A/B testing.
- CloudFront can return different objects to viewers based on the device they're using by checking the `User-Agent` header, which includes information about the devices. For example, CloudFront can return different images based on the screen size of their device. Similarly, the function could consider the value of the `Referer` header and cause CloudFront to return the images to bots that have the lowest available resolution.
- Or you could check cookies for other criteria. For example, on a retail website that sells clothing, if you use cookies to indicate which color a user chose for a jacket, a Lambda function can change the request so that CloudFront returns the image of a jacket in the selected color.
- A Lambda function can generate HTTP responses when CloudFront viewer request or origin request events occur.
- A function can inspect headers or authorization tokens, and insert a header to control access to your content before CloudFront forwards the request to your origin.
- A Lambda function can also make network calls to external resources to confirm user credentials, or fetch additional content to customize a response.

For sample code and additional examples, see [Lambda@Edge Example Functions](#) (p. 367).

Topics

- [Get Started Creating and Using Lambda@Edge Functions](#) (p. 319)
- [Setting IAM Permissions and Roles for Lambda@Edge](#) (p. 329)
- [Writing and Creating a Lambda@Edge Function](#) (p. 335)

- [Adding Triggers for a Lambda@Edge Function \(p. 339\)](#)
- [Testing and Debugging Lambda@Edge Functions \(p. 343\)](#)
- [CloudWatch Metrics and CloudWatch Logs for Lambda Functions \(p. 351\)](#)
- [Deleting Lambda@Edge Functions and Replicas \(p. 352\)](#)
- [Lambda@Edge Event Structure \(p. 353\)](#)
- [Working with Requests and Responses \(p. 363\)](#)
- [Lambda@Edge Example Functions \(p. 367\)](#)
- [Requirements and Restrictions on Lambda Functions \(p. 393\)](#)

Get Started Creating and Using Lambda@Edge Functions

You can use Lambda@Edge functions to do lots of useful things, but it can seem a little complicated when you're getting started. This section explains, at a high level, how Lambda@Edge works with CloudFront and [provides a tutorial](#) that steps through a simple example.

Tip


After you're familiar with how Lambda@Edge works and you've created a Lambda@Edge function, learn more about how you can use Lambda@Edge for your own custom solutions. Learn more about [creating and updating functions](#), [the event structure](#), and [adding CloudFront triggers](#). You can also find more ideas and get code samples in [Lambda@Edge Example Functions \(p. 367\)](#).

Here's an overview of how to create and use Lambda functions with CloudFront:

1. In the AWS Lambda console, create a Lambda function in the US East (N. Virginia) Region. (Or you can create the function programmatically, for example, by using one of the AWS SDKs.)
2. Save and publish a numbered version of the function.

If you want to change the function, you must edit the \$LATEST version of the function in the US East (N. Virginia) Region. Then, before you set it up to work with CloudFront, you publish a new numbered version.

3. Choose the CloudFront distribution and cache behavior that the function applies to. Then specify one or more CloudFront events (*triggers*) that cause the function to execute. For example, you can create a trigger for the function to execute when CloudFront receives a request from a viewer.
4. When you create a trigger, Lambda replicates the function to AWS locations around the world.



Topics

- [Tutorial: Creating a Simple Lambda@Edge Function \(p. 320\)](#)

Tutorial: Creating a Simple Lambda@Edge Function

This tutorial shows you how to get started with Lambda@Edge by helping you create and add a sample Node.js function that runs in CloudFront. The example that we walk through adds HTTP security headers to a response, which can improve security and privacy for a website. You don't need a website for this walkthrough. In it, we simply add security headers to a response when CloudFront retrieves a file.

This example describes the steps to create and configure a Lambda@Edge function. When you create your own Lambda@Edge solution, you follow similar steps and choose from the same options.

Topics

- [Step 1: Sign Up for an AWS Account \(p. 321\)](#)
- [Step 2: Create a CloudFront Distribution \(p. 321\)](#)
- [Step 3: Create Your Function \(p. 321\)](#)

- [Step 4: Add a CloudFront Trigger to Run the Function \(p. 325\)](#)
- [Step 5: Verify that the Function Runs \(p. 327\)](#)
- [Step 6: Troubleshoot Issues \(p. 328\)](#)
- [Step 7: Clean Up Your Example Resources \(p. 328\)](#)
- [Resources for Learning More \(p. 329\)](#)

Step 1: Sign Up for an AWS Account

If you haven't already done so, sign up for Amazon Web Services at <https://aws.amazon.com/>. Choose **Sign Up Now** and enter the required information.

Step 2: Create a CloudFront Distribution

Before you create the example Lambda@Edge function, you must have a CloudFront environment to work with that includes an origin to serve content from.

Are you new to CloudFront? CloudFront delivers content through a worldwide network of edge locations. When you set up a Lambda function with CloudFront, the function can customize content closer to viewers, improving performance. If you're not familiar with CloudFront, take a few minutes before you complete the tutorial to [read a short overview](#) and [learn a bit about how CloudFront caches and serves content](#).

For this example, you create a CloudFront distribution that uses an Amazon S3 bucket as the origin for the distribution. If you already have an environment to use, you can skip this step.

To create a CloudFront distribution with an Amazon S3 origin

1. Create an Amazon S3 bucket with a file or two, such as image files, for sample content. For help, follow the steps in [Upload your content to Amazon S3](#). Make sure that you set permissions to grant public read access to the objects in your bucket.
2. Create a CloudFront distribution and add your S3 bucket as an origin, by following the steps in [Create a CloudFront web distribution](#). If you already have a distribution, you can add the bucket as an origin for that distribution instead.

Tip

Make a note of your distribution ID. Later in this tutorial when you add a CloudFront trigger for your function, you must choose the ID for your distribution in a drop-down list—for example, E653W2221KDDL.

Step 3: Create Your Function

In this step, you create a Lambda function, starting with a blueprint template that's provided in the Lambda Console. The function adds code to update security headers in your CloudFront distribution.

Are you new to Lambda or Lambda@Edge? Lambda@Edge lets you use CloudFront triggers to invoke a Lambda function. When you associate a CloudFront distribution with a Lambda function, CloudFront [intercepts requests and responses](#) at CloudFront edge locations and runs the function. Lambda functions can improve security or customize information close to your viewers, to improve performance. In this tutorial, the function that we create updates the security headers in a CloudFront response.

There are several steps to take when you create a Lambda function. In this tutorial, you use a blueprint template as the basis for your function, and then update the function with code that sets the security headers. Finally, you add and deploy a CloudFront trigger to run the function.

To create a Lambda function

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

Important

Make sure that you're in the **US-East-1 (N. Virginia)** Region (**us-east-1**). You must be in this Region to create Lambda@Edge functions.

2. Choose **Create function**.
3. On the **Create function** page, choose **Use a blueprint**, and then filter for the CloudFront blueprints by entering **cloudfront** in the search field. The **Keyword : cloudfront** is shown, and all the blueprints that are tagged for CloudFront are listed.

Note

CloudFront blueprints are available only in the **US-East-1 (N. Virginia)** Region (**us-east-1**).

4. Choose the **cloudfront-modify-response-header** blueprint as the template for your function.
5. Enter the following information about your function:

Name

Enter a name for your function.

Execution role

Choose how to set the permissions for your function. To use the recommended basic Lambda@Edge permissions policy template, choose **Create a new role from AWS policy templates**.

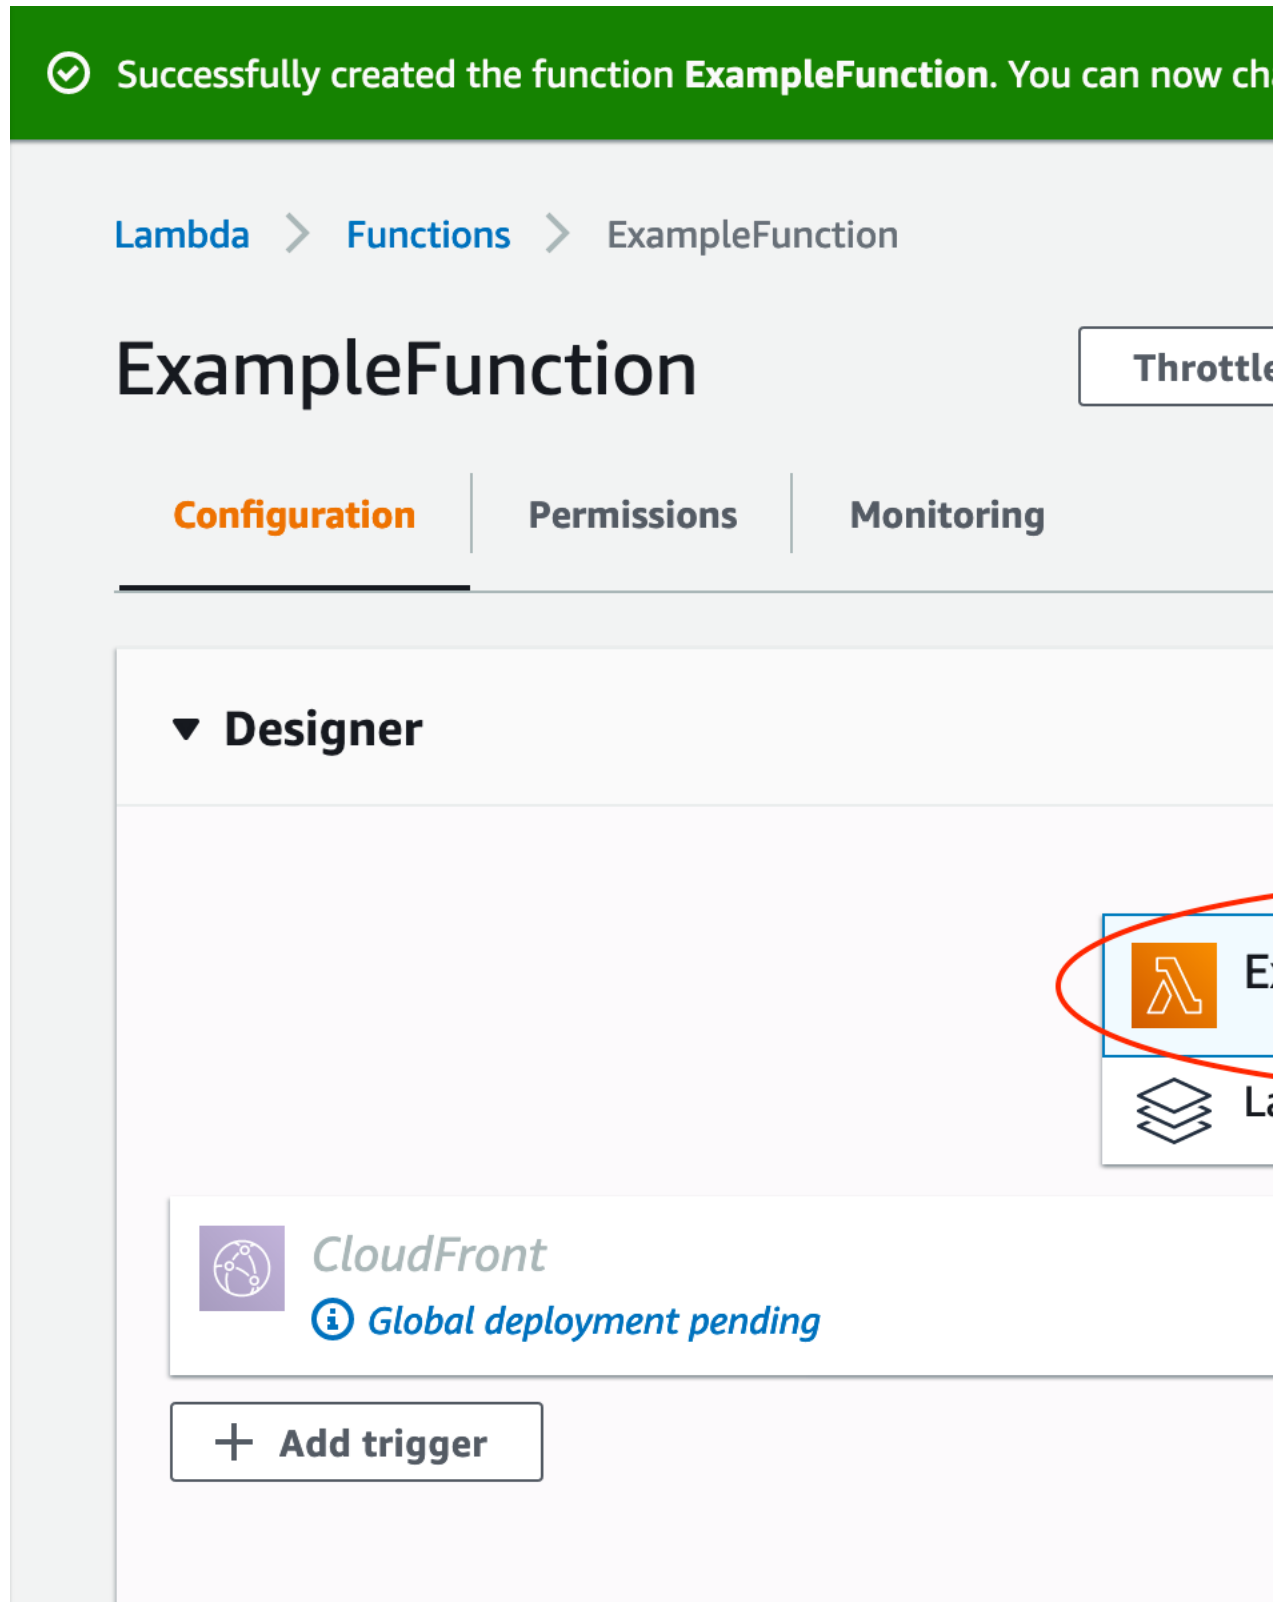
Role name

Enter a name for the role that the policy template creates.

Policy templates

Lambda automatically adds the policy template **Basic Edge Lambda permissions** because you chose a CloudFront blueprint as the basis for your function. This policy template adds execution role permissions that allow CloudFront to run your Lambda function for you in CloudFront locations around the world. For more information, see [Setting IAM Permissions and Roles for Lambda@Edge](#) (p. 329).

6. Choose **Create function**. Lambda creates the function, and on the next page you see your function configuration.
7. In the **Designer** section of the page, choose your function name, as shown in the following image. In this example, the function name is *ExampleFunction*.



8. Scroll down to the **Function code** section of the page, as shown in the following image.

ExampleFunction

Throttle


Function code [Info](#)

Code entry type

Edit code inline ▼

Runtime

Node.js



```
File Edit Find View Go Tools Window
Environment
  ExampleFunction
    index.js
index.js
1 exports.handler =
2   const response =
3   const headers =
4
5   const headers =
6   const headers =
7
8   if (headers[
9     headers[
10       key:
11       valu
12     ]];
13     console.
14
15   }
16
17   return respo
18 };
19
```

Replace the template code with a function that modifies the security headers that your origin returns. For example, you could use code similar to the following:

```
'use strict';
exports.handler = (event, context, callback) => {

    //Get contents of response
    const response = event.Records[0].cf.response;
    const headers = response.headers;

    //Set new headers
    headers['strict-transport-security'] = [{key: 'Strict-Transport-Security', value:
'max-age= 63072000; includeSubdomains; preload'}];
    headers['content-security-policy'] = [{key: 'Content-Security-Policy', value:
"default-src 'none'; img-src 'self'; script-src 'self'; style-src 'self'; object-src
'none'"}];
    headers['x-content-type-options'] = [{key: 'X-Content-Type-Options', value:
'nosniff'}];
    headers['x-frame-options'] = [{key: 'X-Frame-Options', value: 'DENY'}];
    headers['x-xss-protection'] = [{key: 'X-XSS-Protection', value: '1; mode=block'}];
    headers['referrer-policy'] = [{key: 'Referrer-Policy', value: 'same-origin'}];

    //Return modified response
    callback(null, response);
};
```

9. Choose **Save** to save your updated code.

Proceed to the next section to add a CloudFront trigger to run the function.

Step 4: Add a CloudFront Trigger to Run the Function

Now that you have a Lambda function to update security headers, configure the CloudFront trigger to run your function to add the headers in any response that CloudFront receives from the origin for your distribution.

To configure the CloudFront trigger for your function

1. In the **Designer** section of the page, choose **CloudFront**, as shown in the following image.

✔ Successfully updated the function **ExampleFunction**.



Lambda > Functions > ExampleFunction

ExampleFunction



Throttle

Configuration | Permissions | Monitoring

▼ Designer

 **CloudFront**
 *Global deployment pending*

+ Add trigger

 E
 L

2. Scroll down to the **Configure triggers** section of the page, then choose **Deploy to Lambda@Edge**.

3. On the **Deploy to Lambda@Edge** page, under **Configure CloudFront trigger**, enter the following information:

Distribution

The CloudFront distribution ID to associate with your function. In the drop-down list, choose the distribution ID.

Cache behavior

The cache behavior to use with the trigger. For this example, leave the value set to *, which means your distribution's default cache behavior. For more information, see [Cache Behavior Settings \(p. 46\)](#) in the [Values That You Specify When You Create or Update a Distribution \(p. 38\)](#) topic.

CloudFront event

The trigger that specifies when your function runs. We want the security headers function to run whenever CloudFront returns a response from the origin. So in the drop-down list, choose **Origin response**. For more information, see [Adding Triggers for a Lambda@Edge Function \(p. 339\)](#).

4. Under **Confirm deploy to Lambda@Edge**, select the check box to acknowledge that the trigger will be deployed and run your function in all AWS locations.
5. Choose **Deploy** to add the trigger and replicate the function to AWS locations worldwide. Then, if necessary, close the **Deploy to Lambda@Edge** page.
6. Wait for the function to replicate. This typically takes several minutes.

You can check to see if replication is finished by [going to the CloudFront console](#) and viewing your distribution. Wait for the distribution status to change from **In Progress** back to **Deployed**, which means that your function has been replicated. To verify that the function works, follow the steps in the next section.

Step 5: Verify that the Function Runs


Now that you've created your Lambda function and configured a trigger to run it for a CloudFront distribution, check to make sure that the function is accomplishing what you expect it to. In this example, we check the HTTP headers that CloudFront returns, to make sure that the security headers are added.

To verify that your Lambda@Edge function adds security headers

1. In a browser, enter the URL for a file in your S3 bucket. For example, you might use a URL similar to `http://d1111111abcdef8.cloudfront.net/image.jpg`.

For more information about the CloudFront domain name to use in the file URL, see [Customizing the URL Format for Files in CloudFront \(p. 106\)](#).

2. Open your browser's Web Developer toolbar. For example, in your browser window in Chrome, open the context (right-click) menu, and then choose **Inspect**.
3. Choose the **Network** tab.
4. Reload the page to view your image, and then choose an HTTP request on the left pane. You see the HTTP headers displayed in a separate pane.
5. Look through the list of HTTP headers to verify that the expected security headers are included in the list. For example, you might see headers similar to those shown in the following screenshot.



If the security headers are included in your headers list, great! You've successfully created your first Lambda@Edge function. If CloudFront returns errors or there are other issues, continue to the next step to troubleshoot the issues.

Step 6: Troubleshoot Issues

If CloudFront returns errors or doesn't add the security headers as expected, you can investigate your function's execution by looking at CloudWatch Logs. Be sure to use the logs stored in the AWS location that is closest to the location where the function is executed.

For example, if you view the file from London, try changing the Region in the CloudWatch console to Europe (London).

To examine CloudWatch logs for your Lambda@Edge function

1. Sign in to the AWS Management Console and open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. Change **Region** to the location that is shown when you view the file in your browser. This is where the function is executing.
3. In the left pane, choose **Logs** to view the logs for your distribution.

For more information, see [Monitoring CloudFront with Amazon CloudWatch \(p. 428\)](#).

Step 7: Clean Up Your Example Resources

If you created an Amazon S3 bucket and CloudFront distribution just for this tutorial, delete the AWS resources that you allocated so that you no longer accrue charges. After you delete your AWS resources, any content that you added is no longer available.

Tasks

- [Delete the S3 Bucket \(p. 329\)](#)
- [Delete the CloudFront Distribution \(p. 329\)](#)

Delete the S3 Bucket

Before you delete your Amazon S3 bucket, make sure that logging is disabled for the bucket. Otherwise, AWS continues to write logs to your bucket as you delete it.

To disable logging for a bucket

1. Open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. Select your bucket, and then choose **Properties**.
3. From **Properties**, choose **Logging**.
4. Clear the **Enabled** check box.
5. Choose **Save**.

Now, you can delete your bucket. For more information, see [How Do I Delete an S3 Bucket?](#) in the *Amazon Simple Storage Service Console User Guide*.

Delete the CloudFront Distribution

Before you delete a CloudFront distribution, you must disable it. A disabled distribution is no longer functional and does not accrue charges. You can enable a disabled distribution at any time. After you delete a disabled distribution, it's no longer available.

To disable and delete a CloudFront distribution

1. Open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
2. Select the distribution that you want to disable, and then choose **Disable**.
3. When prompted for confirmation, choose **Yes, Disable**.
4. Select the disabled distribution, and then choose **Delete**.
5. When prompted for confirmation, choose **Yes, Delete**.

Resources for Learning More

Now that you have a basic idea of how Lambda@Edge functions work, learn more by reading the following:

- [Lambda@Edge Example Functions \(p. 367\)](#)
- [Lambda@Edge Design Best Practices](#)
- [Reducing Latency and Shifting Compute to the Edge with Lambda@Edge](#)

Setting IAM Permissions and Roles for Lambda@Edge

To configure Lambda@Edge, you must set up specific IAM permissions and an IAM execution role. Lambda@Edge also creates service-linked roles to replicate Lambda functions to CloudFront Regions and to enable CloudWatch to use CloudFront log files.

Topics

- [IAM Permissions Required to Associate Lambda Functions with CloudFront Distributions \(p. 330\)](#)
- [Function Execution Role for Service Principals \(p. 330\)](#)
- [Service-Linked Roles for Lambda@Edge \(p. 331\)](#)

IAM Permissions Required to Associate Lambda Functions with CloudFront Distributions

In addition to the IAM permissions that you need to use AWS Lambda, the IAM user needs the following IAM permissions to associate Lambda functions with CloudFront distributions:

- `lambda:GetFunction`

Allows the user to get configuration information for the Lambda function and a presigned URL to download a .zip file that contains the function.

For the resource, specify the ARN of the function version that you want to execute when a CloudFront event occurs, as shown in the following example:

```
arn:aws:lambda:us-east-1:123456789012:function:TestFunction:2
```

- `lambda:EnableReplication*`

Adds a permission to the resource policy that gives the Lambda replication service permission to get function code and configuration.

Important

The asterisk (*) at the end of the permission is required: `lambda:EnableReplication*`

For the resource, specify the ARN of the function (without the version number) that you want to execute when a CloudFront event occurs, as shown in the following example:

```
arn:aws:lambda:us-east-1:123456789012:function:TestFunction
```

- `iam:CreateServiceLinkedRole`

Allows the user to create a service linked role that is used by Lambda@Edge to replicate Lambda functions in CloudFront. After this role has been created by the first distribution you use with Lambda@Edge, you don't need to add permission to other distributions that you use with Lambda@Edge.

- `cloudfront:UpdateDistribution` or `cloudfront:CreateDistribution`

Use `cloudfront:UpdateDistribution` to update a distribution or `cloudfront:CreateDistribution` to create a distribution.

For more information, see the following documentation:

- [Identity and Access Management in CloudFront \(p. 476\)](#) in this guide.
- [Authentication and Access Control for AWS Lambda](#) in the *AWS Lambda Developer Guide*

Function Execution Role for Service Principals

You must create an IAM role that can be assumed by the service principals `lambda.amazonaws.com` and `edgelambda.amazonaws.com`. This role is assumed by the service principals when they execute

your function. For more information, see [Creating the Roles and Attaching the Policies \(Console\)](#) in the topic "AWS Managed Policies for Job Functions" in the *IAM User Guide*.

You add this role under the **Trust Relationship** tab in IAM (do not add it under the **Permissions** tab).

Here's an example role trust policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": [
          "lambda.amazonaws.com",
          "edgelambda.amazonaws.com"
        ]
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

For information about the permissions that you need to grant to the execution role, see [Manage Permissions: Using an IAM Role \(Execution Role\)](#) in the *AWS Lambda Developer Guide*. Note the following:

- By default, whenever a CloudFront event triggers a Lambda function, data is written to CloudWatch Logs. If you want to use these logs, the execution role needs permission to write data to CloudWatch Logs. You can use the predefined `AWSLambdaBasicExecutionRole` to grant permission to the execution role.

For more information about CloudWatch Logs, see [CloudWatch Metrics and CloudWatch Logs for Lambda Functions](#) (p. 351).

- If your Lambda function code accesses other AWS resources, such as reading an object from an S3 bucket, the execution role needs permission to perform that operation.

Service-Linked Roles for Lambda@Edge

Lambda@Edge uses AWS Identity and Access Management (IAM) [service-linked roles](#). A service-linked role is a unique type of IAM role that is linked directly to a service. Service-linked roles are predefined by the service and include all of the permissions that the service requires to call other AWS services on your behalf.

Lambda@Edge uses the following IAM service-linked role:

- **AWSServiceRoleForLambdaReplicator**—Lambda@Edge uses this role to allow Lambda@Edge to replicate functions to AWS Regions.
- **AWSServiceRoleForCloudFrontLogger**—CloudFront uses this role to push log files into your CloudWatch account, to help you to debug Lambda@Edge validation errors.

When you first add a Lambda@Edge trigger in CloudFront, a role named `AWSServiceRoleForLambdaReplicator` is automatically created to allow Lambda@Edge to replicate functions to AWS Regions. This role is required for using Lambda@Edge functions. The ARN for the `AWSServiceRoleForLambdaReplicator` role looks like this:

```
arn:aws:iam::123456789012:role/aws-service-role/
replicator.lambda.amazonaws.com/AWSServiceRoleForLambdaReplicator
```

The second role, named `AWSServiceRoleForCloudFrontLogger`, is created automatically when you add Lambda@Edge function association to allow CloudFront to push Lambda@Edge error log files to CloudWatch. The ARN for the `AWSServiceRoleForCloudFrontLogger` role looks like this:

```
arn:aws:iam::account_number:role/aws-service-role/  
logger.cloudfront.amazonaws.com/AWSServiceRoleForCloudFrontLogger
```

A service-linked role makes setting up and using Lambda@Edge easier because you don't have to manually add the necessary permissions. Lambda@Edge defines the permissions of its service-linked roles, and only Lambda@Edge can assume the roles. The defined permissions include the trust policy and the permissions policy. The permissions policy cannot be attached to any other IAM entity.

You must remove any associated CloudFront or Lambda@Edge resources before you can delete a service-linked role. This helps protect your Lambda@Edge resources by making sure that you don't remove a service-linked role that is still required to access active resources.

For information about other services that support service-linked roles, see [AWS Services That Work with IAM](#) and look for the services that have **Yes** in the **Service-Linked Role** column.

Service-Linked Role Permissions for Lambda@Edge

Lambda@Edge uses two service-linked roles, named **`AWSServiceRoleForLambdaReplicator`** and **`AWSServiceRoleForCloudFrontLogger`**. The following sections describe the permissions for each of these roles.

Service-Linked Role Permissions for Lambda Replicator

This service-linked role allows Lambda to replicate Lambda@Edge functions to AWS Regions.

The `AWSServiceRoleForLambdaReplicator` service-linked role trusts the following service to assume the role:

- `replicator.lambda.amazonaws.com`

The role permissions policy allows Lambda@Edge to complete the following actions on the specified resources:

- Action: `lambda:CreateFunction` on `arn:aws:lambda:*:*:function:*`
- Action: `lambda>DeleteFunction` on `arn:aws:lambda:*:*:function:*`
- Action: `lambda:DisableReplication` on `arn:aws:lambda:*:*:function:*`
- Action: `iam:PassRole` on all AWS resources
- Action: `cloudfront:ListDistributionsByLambdaFunction` on all AWS resources

Service-Linked Role Permissions for CloudFront Logger

This service-linked role allows CloudFront to push log files into your CloudWatch account, to help you to debug Lambda@Edge validation errors.

The `AWSServiceRoleForCloudFrontLogger` service-linked role trusts the following service to assume the role:

- `logger.cloudfront.amazonaws.com`

The role permissions policy allows Lambda@Edge to complete the following actions on the specified resources:

- Action: `logs:CreateLogGroup` on `arn:aws:logs:*:*:log-group:/aws/cloudfront/*`

- Action: `logs:CreateLogStream` on `arn:aws:logs:*:*:log-group:/aws/cloudfront/*`
- Action: `logs:PutLogEvents` on `arn:aws:logs:*:*:log-group:/aws/cloudfront/*`

You must configure permissions to allow an IAM entity (such as a user, group, or role) to delete the Lambda@Edge service-linked roles. For more information, see [Service-Linked Role Permissions](#) in the *IAM User Guide*.

Creating Service-Linked Roles for Lambda@Edge

You don't typically manually create the service-linked roles for Lambda@Edge. The service creates the roles for you automatically in the following scenarios:

- When you first create a trigger, the service creates a role, `AWSServiceRoleForLambdaReplicator`, if the role doesn't already exist, that allows Lambda to replicate Lambda@Edge functions to AWS Regions.

If you delete the service-linked role, the role will be created again when you add a new trigger for Lambda@Edge in a distribution.

- When you update or create a CloudFront distribution that has a Lambda@Edge association, the service creates a role, `AWSServiceRoleForCloudFrontLogger`, if the role doesn't already exist, that allows CloudFront to push your log files to CloudWatch.

If you delete the service-linked role, the role will be created again when you update or create a CloudFront distribution that has a Lambda@Edge association.

If you must manually create these service-linked roles, run the following commands using the AWS CLI:

To create the `AWSServiceRoleForLambdaReplicator` role

```
aws iam create-service-linked-role --aws-service-name  
replicator.lambda.amazonaws.com
```

To create the `AWSServiceRoleForCloudFrontLogger` role

```
aws iam create-service-linked-role --aws-service-name  
logger.cloudfront.amazonaws.com
```

Editing Lambda@Edge Service-Linked Roles

Lambda@Edge does not allow you to edit the `AWSServiceRoleForLambdaReplicator` or `AWSServiceRoleForCloudFrontLogger` service-linked roles. After the service has created a service-linked role, you cannot change the name of the role because various entities might reference the role. However, you can edit the description of a role by using IAM. For more information, see [Editing a Service-Linked Role](#) in the *IAM User Guide*.

Deleting Lambda@Edge Service-Linked Roles

If you no longer need to use Lambda@Edge, we recommend that you delete the service-linked roles. That way you don't have unused entities that are not actively monitored or maintained. However, you must clean up the Lambda@Edge resources in your account before you can manually delete the roles.

To remove all Lambda@Edge associations from your distributions, update your distributions to remove all Lambda@Edge function triggers, or remove the distributions that use Lambda@Edge functions. For more information, see [Deleting Lambda@Edge Functions and Replicas](#) (p. 352).

After you have removed all Lambda@Edge function associations from your distributions and CloudFront has removed the function replicas from AWS locations, then you can delete the service-linked roles.

Note


If CloudFront hasn't finished updating, service-linked role deletion might fail. If that happens, wait for a few minutes, and then try the deletion steps again.

You must follow separate procedures to manually delete each service-linked role:

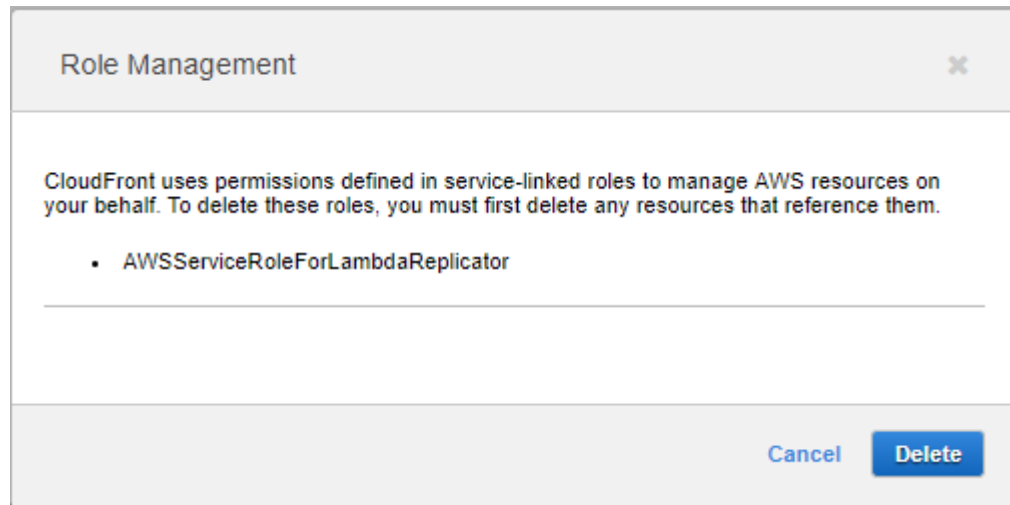
- You delete the `AWSServiceRoleForLambdaReplicator` role by using the CloudFront console.
- You delete the `AWSServiceRoleForCloudFrontLogger` role by using the IAM console.

To manually delete the `AWSServiceRoleForLambdaReplicator` service-linked role

1. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
2. On the **CloudFront Distributions** page, click the avatar in the upper right.



3. Choose **Delete**.



To manually delete the `AWSServiceRoleForCloudFrontLogger` service-linked roles

1. Sign in to the AWS Management Console and open the IAM console at <https://console.aws.amazon.com/iam/>.
2. In the navigation pane of the IAM console, choose **Roles**. Then select the check box next to the role name that you want to delete, not the name or row itself.
3. For **Role** actions at the top of the page, choose **Delete** role.
4. In the confirmation dialog box, review the service last accessed data, which shows when each of the selected roles last accessed an AWS service. This helps you to confirm whether the role is currently active. If you want to proceed, choose **Yes, Delete** to submit the service-linked role for deletion.
5. Watch the IAM console notifications to monitor the progress of the service-linked role deletion. Because the IAM service-linked role deletion is asynchronous, after you submit the role for deletion, the deletion task can succeed or fail. For more information, see [Deleting a Service-Linked Role](#) in the *IAM User Guide*.

Supported Regions for CloudFront Service-Linked Roles

CloudFront supports using service-linked roles for Lambda@Edge in the following Regions:

- US East (Ohio) – `us-east-2`
- US East (N. Virginia) – `us-east-1`
- US West (Oregon) – `us-west-2`
- Asia Pacific (Mumbai) – `ap-south-1`
- Asia Pacific (Seoul) – `ap-northeast-2`
- Asia Pacific (Singapore) – `ap-southeast-1`
- Asia Pacific (Sydney) – `ap-southeast-2`
- Asia Pacific (Tokyo) – `ap-northeast-1`
- Europe (Frankfurt) – `eu-central-1`
- Europe (Ireland) – `eu-west-1`
- Europe (London) – `eu-west-2`
- South America (São Paulo) – `sa-east-1`

Writing and Creating a Lambda@Edge Function

To use Lambda@Edge, you write the code for your Lambda function, then set up AWS Lambda to run the function based on specific CloudFront events (triggers). To set up Lambda to run your function, you use the create function option in Lambda.

You can use the AWS console to work with Lambda functions and CloudFront triggers, or you can work with Lambda@Edge programmatically by using APIs.

- If you use the console, be aware that you can use only the AWS Lambda console to create Lambda functions. You can't use the Amazon CloudFront console to create a function.
- If you want to work with Lambda@Edge programmatically, there are several resources to help you. For more information, see [Creating Lambda Functions and CloudFront Triggers Programmatically \(p. 338\)](#).

Note

You can use either the AWS Lambda console or CloudFront console to add triggers for Lambda@Edge functions.

Topics

- [Writing Functions for Lambda@Edge \(p. 335\)](#)
- [Creating a Lambda@Edge Function in the Lambda Console \(p. 336\)](#)
- [Editing a Lambda Function for Lambda@Edge \(p. 337\)](#)
- [Creating Lambda Functions and CloudFront Triggers Programmatically \(p. 338\)](#)

Writing Functions for Lambda@Edge

There are several resources to help you with writing Lambda@Edge functions:

- For more information about writing Lambda functions that you can use with Lambda@Edge, see [Requirements and Restrictions on Lambda Functions \(p. 393\)](#).
- To learn about the event structure to use with Lambda@Edge functions, see [Lambda@Edge Event Structure \(p. 353\)](#).

- To see examples of Lambda@Edge functions, such as functions for A/B testing and generating an HTTP redirect, see [Lambda@Edge Example Functions \(p. 367\)](#).

The programming model for using Node.js or Python with Lambda@Edge is the same as using Lambda in an AWS Region. For more information, see [Building Lambda Functions with Node.js](#) or [Building Lambda Functions with Python](#).

In your Lambda@Edge code, include the `callback` parameter and return the applicable object for request or response events:

- **Request events** – Include the `cf.request` object in the response.

If you're generating a response, include the `cf.response` object in the response. For more information, see [Generating HTTP Responses in Request Triggers \(p. 364\)](#).

- **Response events** – Include the `cf.response` object in the response.

Creating a Lambda@Edge Function in the Lambda Console

To set up AWS Lambda to run Lambda functions that are based on CloudFront events, follow this procedure.

To create a Lambda@Edge function

1. Sign in to the AWS Management Console and open the AWS Lambda console at <https://console.aws.amazon.com/lambda/>.
2. If you already have one or more Lambda functions, choose **Create function**.

If you've don't have any functions, choose **Get Started Now**.
3. In the Region list at the top of the page, choose **US East (N. Virginia)**.
4. Create a function using your own code or create a function starting with a CloudFront blueprint.
 - To create a function using your own code, choose **Author from scratch**.
 - To display a list of blueprints for CloudFront, type **cloudfront** in the filter field, and then press **Enter**.

If you find a blueprint that you want to use, choose the name of the blueprint.

5. In the **Basic information** section, specify the following values:

Name

Type a name for your function.

Role

Choose **Create new role from template(s)**.

Note

Choosing this value will get you started quickly. Or you can choose **Choose an existing role** or **Create a custom role**. If you choose one of these, follow the prompts to complete the information for this section.

Role name

Type a name for the role.

Policy templates

Choose **Basic Edge Lambda permissions**.

6. If you chose **Author from scratch** in step 4, skip to step 7.

If you chose a blueprint in step 4, the **cloudfront** section lets you create one trigger, which associates this function with a cache in a CloudFront distribution and a CloudFront event. We recommend that you choose **Remove** at this point, so there isn't a trigger for the function when it's created. Then you can add triggers later.

Important

Why add triggers later? Generally it's best to test and debug the function before you add triggers. If you choose instead to add a trigger now, the function will start to run as soon as you create the function and it finishes replicating to AWS locations around the world, and the corresponding distribution is deployed.

7. Choose **Create function**.

Lambda creates two versions of your function: **\$LATEST** and **Version 1**. You can edit only the **\$LATEST** version, but the console initially displays **Version 1**.

8. To edit the function, choose **Version 1** near the top of the page, under the ARN for the function. Then, on the **Versions** tab, choose **\$LATEST**. (If you left the function and then returned to it, the button label is **Qualifiers**.)
9. On the **Configuration** tab, choose the applicable **Code entry type**. Then follow the prompts to edit or upload your code.
10. For **Runtime**, choose the value based on your function's code.
11. In the **Tags** section, add any applicable tags.
12. Choose **Actions**, and then choose **Publish new version**.
13. Type a description for the new version of the function.
14. Choose **Publish**.
15. Test and debug the function. For more information about testing in the Lambda console, see the *Invoke the Lambda Function and Verify Results, Logs, and Metrics* section in [Create a Lambda Function with the Console](#) in the *AWS Lambda Developer Guide*.
16. When you're ready to have the function execute for CloudFront events, publish another version and edit the function to add triggers. For more information, see [Adding Triggers for a Lambda@Edge Function](#) (p. 339).

Editing a Lambda Function for Lambda@Edge

When you want to edit a Lambda function, note the following:

- The original version is labeled **\$LATEST**.
- You can edit only the **\$LATEST** version.
- Each time you edit the **\$LATEST** version, you must publish a new numbered version.
- You can't create triggers for **\$LATEST**.
- When you publish a new version of a function, Lambda doesn't automatically copy triggers from the previous version to the new version. You must reproduce the triggers for the new version.
- When you add a trigger for a CloudFront event to a function, if there's already a trigger for the same distribution, cache behavior, and event for an earlier version of the same function, Lambda deletes the trigger from the earlier version.
- After you make updates to a CloudFront distribution, like adding triggers, you must wait for the changes to propagate to edge locations before the functions you've specified in the triggers will work.

To edit a Lambda function (AWS Lambda console)

1. Sign in to the AWS Management Console and open the AWS Lambda console at <https://console.aws.amazon.com/lambda/>.
2. In the Region list at the top of the page, choose **US East (N. Virginia)**.
3. In the list of functions, choose the name of the function that you want to edit.

By default, the console displays the **\$LATEST** version. You can view earlier versions (choose **Qualifiers**), but you can only edit **\$LATEST**.
4. On the **Code** tab, for **Code entry type**, choose to edit the code in the browser, upload a .zip file, or upload a file from Amazon S3.
5. Choose either **Save** or **Save and test**.
6. Choose **Actions**, and choose **Publish new version**.
7. In the **Publish new version from \$LATEST** dialog box, enter a description of the new version. This description appears in the list of versions, along with an automatically generated version number.
8. Choose **Publish**.

The new version automatically becomes the latest version. The version number appears on the **Version** button in the upper-left corner of the page.

9. Choose the **Triggers** tab.
10. Choose **Add trigger**.
11. In the **Add trigger** dialog box, choose the dotted box, and then choose **CloudFront**.

Note

If you've already created one or more triggers for a function, CloudFront is the default service.

12. Specify the following values to indicate when you want the Lambda function to execute.

Distribution ID

Choose the ID of the distribution that you want to add the trigger to.

Cache behavior

Choose the cache behavior that specifies the objects that you want to execute the function on.

CloudFront event

Choose the CloudFront event that causes the function to execute.

Enable trigger and replicate

Select this check box so Lambda replicates the function to Regions globally.

13. Choose **Submit**.
14. To add more triggers for this function, repeat steps 10 through 13.

Creating Lambda Functions and CloudFront Triggers Programmatically

You can set up Lambda@Edge functions and CloudFront triggers programmatically by using API actions instead of by using the AWS console. For more information, see the following:

- [API Reference](#) in the *AWS Lambda Developer Guide*
- [Amazon CloudFront API Reference](#)

- **AWS CLI**
 - [Lambda create-function command](#)
 - [CloudFront create-distribution command](#)
 - [CloudFront create-distribution-with-tags command](#)
 - [CloudFront update-distribution command](#)
- [AWS SDKs](#) (See the **SDKs & Toolkits** section.)
- [AWS Tools for PowerShell Cmdlet Reference](#)

Adding Triggers for a Lambda@Edge Function

A Lambda@Edge trigger is one combination of CloudFront distribution, cache behavior, and event that causes a function to execute. You can specify one or more CloudFront triggers that cause the function to run. For example, you can create a trigger that causes the function to execute when CloudFront receives a request from a viewer for a specific cache behavior you set up for your distribution.

Tip

If you're not familiar with CloudFront cache behaviors, here's a brief overview. When you create a CloudFront distribution, you specify settings that tell CloudFront how to respond when it receives different requests. The default settings are called the default cache behavior for the distribution. You can set up additional cache behaviors that define how CloudFront responds under specific circumstances, for example, when it receives a request for a specific file type. For more information, see [Cache Behavior Settings](#).

At the time that you create a Lambda function, you can specify only one trigger. But you can add more triggers to the same function later in one of two ways: by using the Lambda console or by editing the distribution in the CloudFront console.

- Using the Lambda console works well if you want to add more triggers to a function for the same CloudFront distribution.
- Using the CloudFront console can be better if you want to add triggers for multiple distributions because it's easier to find the distribution that you want to update. You can also update other CloudFront settings at the same time.

Note


If you want to work with Lambda@Edge programmatically, there are several resources to help you. For more information, see [Creating Lambda Functions and CloudFront Triggers Programmatically](#) (p. 338).

Topics

- [CloudFront Events That Can Trigger a Lambda Function](#) (p. 339)
- [How to Decide Which CloudFront Event to Use to Trigger a Lambda Function](#) (p. 341)
- [Adding Triggers by Using the Lambda Console](#) (p. 341)
- [Adding Triggers by Using the CloudFront Console](#) (p. 342)

CloudFront Events That Can Trigger a Lambda Function

For each cache behavior in a CloudFront distribution, you can add up to four triggers (associations) that cause a Lambda function to execute when specific CloudFront events occur. CloudFront triggers can be based on one of four CloudFront events, as shown in the following diagram.



The CloudFront events that can be used to trigger Lambda@Edge functions are the following:

Viewer Request

The function executes when CloudFront receives a request from a viewer, before it checks to see whether the requested object is in the CloudFront cache.

Origin Request

The function executes *only* when CloudFront forwards a request to your origin. When the requested object is in the CloudFront cache, the function doesn't execute.

Origin Response

The function executes after CloudFront receives a response from the origin and before it caches the object in the response. Note that the function executes even if an error is returned from the origin.

The function doesn't execute in the following cases:

- When the requested file is in the CloudFront cache and is not expired.
- When the response is generated from a function that was triggered by an origin request event.

Viewer Response

The function executes before returning the requested file to the viewer. Note that the function executes regardless of whether the file is already in the CloudFront cache.

The function doesn't execute in the following cases:

- When the origin returns an HTTP status code of 400 or higher.
- When a custom error page is returned.
- When the response is generated from a function that was triggered by a viewer request event.
- When CloudFront automatically redirects an HTTP request to HTTPS (when the value of [Viewer Protocol Policy](#) (p. 49) is **Redirect HTTP to HTTPS**).

When you add multiple triggers to the same cache behavior, you can use them to run the same function or run different functions for each trigger. You can also associate the same function with more than one distribution.

Note

When a CloudFront event triggers the execution of a Lambda function, the function must finish before CloudFront can continue. For example, if a Lambda function is triggered by a CloudFront viewer request event, CloudFront won't return a response to the viewer or forward the request to the origin until the Lambda function finishes running. This means that each request that triggers a Lambda function increases latency for the request, so you'll want the function to execute as fast as possible.

How to Decide Which CloudFront Event to Use to Trigger a Lambda Function

When you're deciding which CloudFront event you want to use to trigger a Lambda function, consider the following:

Do you want CloudFront to cache objects that are changed by a Lambda function?

If you want CloudFront to cache an object that was modified by a Lambda function so that CloudFront can serve the object from the edge location the next time it's requested, use the origin request or origin response event. This reduces the load on the origin, reduces latency for subsequent requests, and reduces the cost of invoking Lambda@Edge on subsequent requests.

For example, if you want to add, remove, or change headers for objects that are returned by the origin and you want CloudFront to cache the result, use the origin response event.

Do you want the function to execute for every request?

If you want the function to execute for every request that CloudFront receives for the distribution, use the viewer request or viewer response events. Origin request and origin response events occur only when a requested object isn't cached in an edge location and CloudFront forwards a request to the origin.

Does the function change the cache key?

If you want the function to change a value that you're using as a basis for caching, use the viewer request event. For example, if a function changes the URL to include a language abbreviation in the path (for example, because the user chose their language from a dropdown list), use the viewer request event:

- **URL in the viewer request** – `http://example.com/en/index.html`
- **URL when the request comes from an IP address in Germany** – `http://example.com/de/index.html`

You also use the viewer request event if you're caching based on cookies or request headers.

Note

If the function changes cookies or headers, configure CloudFront to forward the applicable part of the request to the origin. For more information, see the following topics:

- [Caching Content Based on Cookies \(p. 244\)](#)
- [Caching Content Based on Request Headers \(p. 246\)](#)

Does the function affect the response from the origin?

If you want the function to change the request in a way that affects the response from the origin, use the origin request event. Typically, most viewer request events aren't forwarded to the origin; CloudFront responds to a request with an object that's already in the edge cache. If the function changes the request based on an origin request event, CloudFront caches the response to the changed origin request.

Adding Triggers by Using the Lambda Console

To add triggers to a Lambda@Edge function (AWS Lambda console)

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

2. In the Region list at the top of the page, choose **US East (N. Virginia)**.
3. On the **Functions** page, choose the name of the function that you want to add triggers for.
4. Choose **Qualifiers**, and then choose the **Versions** tab.
5. Choose the version that you want to add triggers to.

Important

You can't create triggers for the `$LATEST` version, you must create them for a numbered version.

After you choose a version, the name of the button changes to **Version: `$LATEST`** or **Version: *version number***.

6. Choose the **Triggers** tab.
7. Choose **Add triggers**.
8. In the **Add trigger** dialog box, choose the dotted box, and then choose **CloudFront**.

Note

If you've already created one or more triggers, CloudFront is the default service.

9. Specify the following values to indicate when you want the Lambda function to execute.

Distribution ID

Choose the ID of the distribution that you want to add the trigger to.

Cache behavior

Choose the cache behavior that specifies the objects that you want to execute the function on.

Note

If you specify `*` for the cache behavior, the Lambda function deploys to the default cache behavior.

CloudFront event

Choose the CloudFront event that causes the function to execute.

Include body

Select this check box if you want to access the request body in your function.

Enable trigger and replicate

Select this check box so that AWS Lambda replicates the function to Regions globally.

10. Choose **Submit**.

The function starts to process requests for the specified CloudFront events when the updated CloudFront distribution is deployed. To determine whether a distribution is deployed, choose **Distributions** in the navigation pane. When a distribution is deployed, the value of the **Status** column for the distribution changes from **In Progress** to **Deployed**.

Adding Triggers by Using the CloudFront Console

To add triggers for CloudFront events to a Lambda function (CloudFront console)

1. Get the ARN of the Lambda function that you want to add triggers for:
 - a. Sign in to the AWS Management Console and open the AWS Lambda console at <https://console.aws.amazon.com/lambda/>.
 - b. In the list of Regions at the top of the page, choose **US East (N. Virginia)**.
 - c. In the list of functions, choose name of the function that you want to add triggers to.

- d. Choose **Qualifiers**, choose the **Versions** tab, and choose the numbered version that you want to add triggers to.

Important

You can add triggers only to a numbered version, not to **\$LATEST**.

- e. Copy the ARN that appears at the top of the page, for example:

```
arn:aws:lambda:us-east-1:123456789012:function:TestFunction:2
```

The number at the end (**2** in this example) is the version number of the function.

2. Open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
3. In the list of distributions, choose the ID of the distribution that you want to add triggers to.
4. Choose the **Behaviors** tab.
5. Select the check box for the cache behavior that you want to add triggers to, and then choose **Edit**.
6. At **Lambda Function Associations**, in the **Event Type** list, choose when you want the function to execute: for viewer requests, viewer responses, origin requests, or origin responses.

For more information, see [How to Decide Which CloudFront Event to Use to Trigger a Lambda Function](#) (p. 341).

7. Paste the ARN of the Lambda function that you want to execute when the chosen event occurs. This is the value that you copied in step 1.
8. Select **Include Body** if you want to access the request body in your function.

If you just want to replace the request body, you don't need to select this option.

9. To execute the same function for more event types, choose **+** and repeat steps 6 and 7.
10. Choose **Yes, Edit**.
11. To add triggers to more cache behaviors for this distribution, repeat steps 5 through 9.

The function starts to process requests for the specified CloudFront events when the updated CloudFront distribution is deployed. To determine whether a distribution is deployed, choose **Distributions** in the navigation pane. When a distribution is deployed, the value of the **Status** column for the distribution changes from **In Progress** to **Deployed**.

Testing and Debugging Lambda@Edge Functions

This topic includes sections that describe strategies for testing and debugging Lambda@Edge functions. It's important to test your Lambda@Edge function code standalone, to make sure that it completes the intended task, and to do integration testing, to make sure that the function works correctly with CloudFront.

During integration testing or after your function has been deployed, you might need to debug CloudFront errors, such as HTTP 5xx errors. Errors can be an invalid response returned from the Lambda function, execution errors when the function is triggered, or errors due to execution throttling by the Lambda service. Sections in this topic share strategies for determining which type of failure is the issue, and then steps you can take to correct the problem.

Note

When you review CloudWatch log files or metrics when you're troubleshooting errors, be aware that they are displayed or stored in the Region closest to the location where the function executed. So, if you have a website or web application with users in the United Kingdom, and you have a Lambda function associated with your distribution, for example, you must change the Region to view the CloudWatch metrics or log files for the London AWS Region. For more information, see *Determining the Lambda@Edge Region* later in this topic.

Topics

- [Testing your Lambda@Edge Functions \(p. 344\)](#)
- [Identifying Lambda Function Errors in CloudFront \(p. 344\)](#)
- [Troubleshooting Invalid Lambda Function Responses \(Validation Errors\) \(p. 349\)](#)
- [Troubleshooting Lambda Function Execution Errors \(p. 350\)](#)
- [Determining the Lambda@Edge Region \(p. 350\)](#)
- [Determining if Your Account Pushes Logs to CloudWatch \(p. 351\)](#)

Testing your Lambda@Edge Functions

There are two steps to testing your Lambda function: standalone testing and integration testing.

Test standalone functionality

Before you add your Lambda function to CloudFront, make sure to test the functionality first by using the testing capabilities in the Lambda console or by using other methods. For more information about testing in the Lambda console, see the *Invoke the Lambda Function and Verify Results, Logs, and Metrics* section in [Create a Lambda Function with the Console](#) in the *AWS Lambda Developer Guide*.

Test your function's operation in CloudFront

It's important to complete integration testing, where your function is associated with a distribution and runs based on a CloudFront event. Make sure that the function is triggered for the right event, and returns a response that is valid and correct for CloudFront. For example, make sure that the event structure correct, that only valid headers are included, and so on.

As you iterate on integration testing with your function in the Lambda console, refer to the steps in the Lambda@Edge tutorial as you modify your code or change the CloudFront trigger that calls your function. For example, make sure that you're working in a numbered version of your function, as described in this step of the tutorial: [Step 4: Add a CloudFront Trigger to Run the Function \(p. 325\)](#).

As you make changes and deploy them, be aware that it will take several minutes for your updated function and CloudFront triggers to replicate across all Regions. This typically takes a few minutes but can take up to 15 minutes.

You can check to see if replication is finished by going to the CloudFront console and viewing your distribution:

- Go to the CloudFront console at <https://console.aws.amazon.com/cloudfront>.

Check for the distribution status to change from **In Progress** back to **Deployed**, which means that your function has been replicated. Then follow the steps in the next section to verify that the function works.

Be aware that testing in the console only validates your function's logic, and does not apply any service quotas (formerly known as limits) that are specific to Lambda@Edge.

Identifying Lambda Function Errors in CloudFront

After you've verified that your function logic works correctly, you might still see HTTP 5xx errors when your function runs in CloudFront. HTTP 5xx errors can be returned for a variety of reasons, which can include Lambda function errors or other issues in CloudFront.

- If you use Lambda@Edge functions, you can use graphs in the CloudFront console to help track down what's causing the error, and then work to fix it. For example, you can see if HTTP 5xx errors are caused by CloudFront or by Lambda functions, and then, for specific functions, you can view related log files to investigate the issue.
- To troubleshoot HTTP errors in general in CloudFront, see the troubleshooting steps in the following topic: [Troubleshooting Error Responses from Your Origin \(p. 254\)](#).

What Causes Lambda Function Errors in CloudFront

There are several reasons why a Lambda function might cause an HTTP 5xx error, and the troubleshooting steps you should take depend on the type of error. Errors can be categorized as the following:

A Lambda function execution error

An execution error results when CloudFront doesn't get a response from Lambda because there are unhandled exceptions in the function or there's an error in the code. For example, if the code includes `callback(Error)`. For more information, see [Lambda Function Errors](#) in the AWS Lambda Developer Guide.

An invalid Lambda function response is returned to CloudFront

After the function runs, CloudFront receives a response from Lambda. An error is returned if the object structure of the response doesn't conform to the [Lambda@Edge Event Structure \(p. 353\)](#), or the response contains invalid headers or other invalid fields.

The execution in CloudFront is throttled due to Lambda service quotas (formerly known as limits)

The Lambda service throttles executions in each Region, and returns an error if you exceed the quota.


How to Determine the Type of Failure

To help you decide where to focus as you debug and work to resolve errors returned by CloudFront, it's helpful to identify why CloudFront is returning an HTTP error. To get started, you can use the graphs provided in the **Monitoring** section of the CloudFront console on the AWS Management Console. For more information about viewing graphs in the **Monitoring** section of the CloudFront console, see [Monitoring CloudFront with Amazon CloudWatch \(p. 428\)](#).

The following graphs can be especially helpful when you want to track down whether errors are being returned by origins or by a Lambda function, and to narrow down the type of issue when it's an error from a Lambda function.

Error rates graph

One of the graphs that you can view on the **Overview** tab for each of your distributions is an **Error rates** graph. This graph displays the rate of errors as a percentage of total requests coming to your distribution. The graph shows the total error rate, total 4xx errors, total 5xx errors, and total 5xx errors from Lambda functions. Based on the error type and volume, you can take steps to investigate and troubleshoot the cause.



- If you see Lambda errors, you can investigate further by looking at the specific types of errors that the function returns. The **Lambda@Edge errors** tab includes graphs that categorize function errors by type to help you pinpoint the issue for a specific function.
- If you see CloudFront errors, you can troubleshoot and work to fix origin errors or change your CloudFront configuration. For more information, see [Troubleshooting Error Responses from Your Origin](#) (p. 254).

Execution errors and invalid function responses graphs


The **Lambda@Edge errors** tab includes graphs that categorize the Lambda@Edge errors for a specific distribution, by type. For example, one graph shows all execution errors by AWS Region. To make it easier to troubleshoot issues, on the same page, you can look for specific problems by opening and examining the log files for specific functions by Region. Under **View execution error logs** or **View invalid function response logs**, choose a Region (and, for execution errors, a function), and then choose **View logs**.

Overview


Lambda@Edge errors

The metrics for Lambda@Edge functions are aggregated in AWS Regions:

Execution errors



Invalid function responses




In addition, read the following sections in this chapter for more recommendations about troubleshooting and fixing errors.

Throttles graph

The **Lambda@Edge errors** tab also includes a **Throttles** graph. On occasion, the Lambda service throttles your function invocations on per Region basis, if you reach the regional concurrency quota (formerly known as limit). If you see a limit exceeded error, your function has reached a quota that the Lambda service imposes on executions in a Region. For more information, including how to request an increase in the quota, see [Quotas on Lambda@Edge \(p. 500\)](#).

Throttles



For an example about how to use this information in troubleshooting HTTP errors, see [Four steps for debugging your content delivery on AWS](#).

Troubleshooting Invalid Lambda Function Responses (Validation Errors)

If you identify that your problem is a Lambda validation error, it means that your Lambda function is returning an invalid response to CloudFront. Follow the guidance in this section to take steps to review your function and make sure that your response conforms to CloudFront requirements.

CloudFront validates the response from a Lambda function in two ways:

- **The Lambda response must conform to the required object structure.** Examples of bad object structure include the following: unparsable JSON, missing required fields, and an invalid object in the response. For more information, see the [Lambda@Edge Event Structure \(p. 353\)](#).
- **The response must include only valid object values.** An error will occur if the response includes a valid object but has values that are not supported. Examples include the following: adding or updating blacklisted or read-only headers (see [Headers \(p. 394\)](#) in the *Requirements and Restrictions on Lambda Functions* topic), exceeding the maximum body size (see *Restrictions on the Size of the Generated Response* in the [Lambda Response Errors \(p. 366\)](#) topic), and invalid characters or values (see the [Lambda@Edge Event Structure \(p. 353\)](#)).

When Lambda returns an invalid response to CloudFront, error messages are written to log files which CloudFront pushes to CloudWatch in the Region of where the Lambda function executed. It's the default behavior to send the log files to CloudWatch when there's an invalid response. However, if you associated a Lambda function with CloudFront before the functionality was released, it might not be enabled for your function. For more information, see *Determine if Your Account Pushes Logs to CloudWatch* later in the topic.

CloudFront pushes log files to the Region corresponding to where your function executed, in the log group that's associated with your distribution. Log groups have the following format: `/aws/cloudfront/LambdaEdge/DistributionId`, where *DistributionId* is your distribution's ID. To determine the Region where you can find the CloudWatch log files, see *Determining the Lambda@Edge Region* later in this topic.

If the error is reproducible, you can create a new request that results in the error and then find the request id in a failed CloudFront response (`x-amz-cf-id` header) to locate a single failure in log files. The log file entry includes information that can help you identify why the error is being returned, and also lists the corresponding Lambda request id so you can analyze the root cause in the context of a single request.

If an error is intermittent, you can use CloudFront access logs to find the request id for a request that has failed, and then search CloudWatch logs for the corresponding error messages. For more information, see the previous section, *Determining the Type of Failure*.

Troubleshooting Lambda Function Execution Errors

If the problem is a Lambda execution error, it can be helpful to create logging statements for Lambda functions, to write messages to CloudWatch log files that monitor the execution of your function in CloudFront and determine if it's working as expected. Then you can search for those statements in the CloudWatch log files to verify that your function is working.

Note

Even if you haven't changed your Lambda@Edge function, updates to the Lambda function execution environment might affect it and could return an execution error. For information about testing and migrating to a later version, see [Upcoming updates to the AWS Lambda and AWS Lambda@Edge execution environment](#).

Determining the Lambda@Edge Region

To see the Regions where your Lambda@Edge function is receiving traffic, view metrics for the function on the CloudFront console on the AWS Management Console. Metrics are displayed for each AWS Region. On the same page, you can choose a Region and view log files for that Region so you can investigate issues. You must review CloudWatch log files in the correct AWS Region to see the log files created when CloudFront executed your Lambda function.

For more information about viewing graphs in the **Monitoring** section of the CloudFront console, see [Monitoring CloudFront with Amazon CloudWatch \(p. 428\)](#).

Determining if Your Account Pushes Logs to CloudWatch

By default, CloudFront enables logging invalid Lambda function responses, and pushes the log files to CloudWatch by using one of the [Service-Linked Roles for Lambda@Edge](#) (p. 331). If you have Lambda@Edge functions that you added to CloudFront before the invalid Lambda function response log feature was released, logging is enabled when you next update your Lambda@Edge configuration, for example, by adding a CloudFront trigger.

You can verify that pushing the log files to CloudWatch is enabled for your account by doing the following:

- **Check to see if the logs appear in CloudWatch.** Make sure that you look in the Region where the Lambda@Edge function executed. For more information, see [Determining the Lambda@Edge Region](#) (p. 350).
- **Determine if the related service-linked role exists in your account in IAM.** To do this, open the IAM console at <https://console.aws.amazon.com/iam/>, and then choose **Roles** to view the list of service-linked roles for your account. Look for the following role: `AWSServiceRoleForCloudFrontLogger`.

CloudWatch Metrics and CloudWatch Logs for Lambda Functions

You can use CloudWatch metrics to monitor, in real time, problems with your Lambda@Edge functions. You can also use CloudWatch Logs to get the function logs. There's no additional charge for metrics or logs.

Topics

- [CloudWatch Metrics](#) (p. 351)
- [CloudWatch Logs](#) (p. 351)

CloudWatch Metrics

When you create a trigger for a CloudFront event, Lambda begins to send metrics to CloudWatch automatically. Metrics are available for all Lambda Regions, but to view metrics in the CloudWatch console or get the metric data from the CloudWatch API, you must use the US East (N. Virginia) Region (us-east-1). The metric group name is formatted as: `AWS/CloudFront/distribution-ID`, where *distribution-ID* is the ID of the CloudFront distribution that the Lambda@Edge function is associated with.

For more information about CloudWatch metrics, see the [Amazon CloudWatch User Guide](#).

CloudWatch Logs

Lambda automatically sends function logs to CloudWatch Logs. You can access the log files using the CloudWatch console or the CloudWatch Logs API.

Lambda creates CloudWatch Logs log streams in the AWS Regions closest to the location where the function is executed. The log group name is formatted as: `/aws/lambda/us-east-1.function-name`, where *function-name* is the name that you gave to the function when you created it.

Note

Lambda@Edge throttles logs based on the request volume and the size of logs.

You must review CloudWatch log files in the correct AWS Region to see your Lambda@Edge function log files. To see the Regions where your Lambda@Edge function is receiving traffic, view graphs of metrics for the function on the CloudFront console. Metrics are displayed for each AWS Region. On the same page, you can choose a Region and then view log files for that Region so that you can investigate issues.

To learn more about how to use CloudWatch Logs with Lambda@Edge functions, see the following:

- For more information about viewing graphs in the **Monitoring** section of the CloudFront console, see [Monitoring CloudFront with Amazon CloudWatch \(p. 428\)](#).
- For information about the permissions required to send data to CloudWatch Logs, see [Setting IAM Permissions and Roles for Lambda@Edge](#) in the *IAM User Guide*.
- For information about adding logging to a Lambda function, see [Function Logging in Node.js](#) or [Function Logging in Python](#) in the *AWS Lambda Developer Guide*.
- For information about CloudWatch Logs quotas (formerly known as limits), see [CloudWatch Logs quotas](#) in the *Amazon CloudWatch Logs User Guide*.

Deleting Lambda@Edge Functions and Replicas

You can delete a Lambda@Edge function only when the replicas of the function have been deleted by CloudFront. Replicas of a Lambda function are automatically deleted in the following situations:

- After you remove the last association for the function from all of your CloudFront distributions. If more than one distribution uses a function, the replicas are deleted only after you remove the function association from the last distribution.
- After you delete the last distribution that a function was associated with.

Replicas are typically deleted within a few hours. You cannot manually delete Lambda@Edge function replicas. This helps prevent a situation where a replica is deleted that is still in use, which would result in an error.

Don't build applications that use Lambda@Edge function replicas outside of CloudFront. These replicas are deleted when their associations with distributions are removed, or when distributions themselves are deleted. The replica that an outside application depends on might be removed without warning, causing it to fail.

To delete a Lambda@Edge function association from a CloudFront distribution (console)

1. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
2. Choose the ID of the distribution that has the Lambda@Edge function association that you want to delete.
3. Choose the **Behaviors** tab.
4. Choose the check box next to the cache behavior that has the Lambda@Edge function association that you want to delete, and then choose **Edit**.
5. Scroll down to **Lambda Function Associations**, and then choose the **X** icon next to each Lambda@Edge function association that you want to delete.
6. Choose **Yes, Edit** to save your changes.

After you delete a Lambda@Edge function association from a CloudFront distribution, you can optionally delete the Lambda function or function version from AWS Lambda. You can also delete a specific version of a Lambda function if the version doesn't have any CloudFront distributions associated with it. If you remove all the associations for a Lambda function version, you can typically delete the function version a few hours later.

Lambda@Edge Event Structure

The following topics describe the request and response event objects that CloudFront passes to a Lambda@Edge function when it's triggered.

Topics

- [Dynamic Origin Selection \(p. 353\)](#)
- [Request Events \(p. 353\)](#)
- [Response Events \(p. 358\)](#)

Dynamic Origin Selection

You can use [the path pattern in a cache behavior \(p. 47\)](#) to route requests to an origin based on the path and name of the requested object, such as `images/*.jpg`. Using Lambda@Edge, you can also route requests to an origin based on other characteristics, such as the values in request headers.

There are a number of ways that this dynamic origin selection can be useful. For example, you can distribute requests across origins in different geographic areas to help with global load balancing. Or you can selectively route requests to different origins that each serve a particular function: bot handling, SEO optimization, authentication, and so on. For code examples that demonstrate how to use this feature, see [Content-Based Dynamic Origin Selection - Examples \(p. 381\)](#).

In the CloudFront origin request event, the `origin` object in the event structure contains information about the origin that the request would be routed to, based on the path pattern. You can update the values in the `origin` object to route a request to a different origin. When you update the `origin` object, you don't need to define the origin in the distribution. You can also replace an Amazon S3 origin object with a custom origin object, and vice versa. You can only specify a single origin per request, though; either a custom origin or an Amazon S3 origin, but not both.

Request Events

The following topics show the structure of the object that CloudFront passes to a Lambda function for [viewer and origin request events \(p. 339\)](#). These examples show a `GET` request with no body. Following the examples is a list of all the possible fields in viewer and origin request events.

Topics

- [Example Viewer Request \(p. 353\)](#)
- [Example Origin Request \(p. 354\)](#)
- [Request Event Fields \(p. 355\)](#)

Example Viewer Request

The following example shows a viewer request event object.

```
{
  "Records": [
    {
      "cf": {
        "config": {
          "distributionDomainName": "d111111abcdef8.cloudfront.net",
          "distributionId": "EDFDVBD6EXAMPLE",
          "eventType": "viewer-request",
          "requestId": "4TyzHTaYWb1GX1qTfsHhEqV6HUUd_BzoBZnwfnvQc_1oF26ClkoUSEQ=="
        },

```

```
    "request": {
      "clientIp": "203.0.113.178",
      "headers": {
        "host": [
          {
            "key": "Host",
            "value": "d1111111abcdef8.cloudfront.net"
          }
        ],
        "user-agent": [
          {
            "key": "User-Agent",
            "value": "curl/7.66.0"
          }
        ],
        "accept": [
          {
            "key": "accept",
            "value": "*/*"
          }
        ]
      },
      "method": "GET",
      "querystring": "",
      "uri": "/"
    }
  }
}
```

Example Origin Request

The following example shows an origin request event object.

```
{
  "Records": [
    {
      "cf": {
        "config": {
          "distributionDomainName": "d1111111abcdef8.cloudfront.net",
          "distributionId": "EDFDVBD6EXAMPLE",
          "eventType": "origin-request",
          "requestId": "4TyzHTaYWb1GX1qTfsHhEqV6HUDd_BzoBZnwfnvQc_1oF26ClkoUSEQ=="
        },
        "request": {
          "clientIp": "203.0.113.178",
          "headers": {
            "x-forwarded-for": [
              {
                "key": "X-Forwarded-For",
                "value": "203.0.113.178"
              }
            ],
            "user-agent": [
              {
                "key": "User-Agent",
                "value": "Amazon CloudFront"
              }
            ],
            "via": [
              {
                "key": "Via",
                "value": "2.0 2afae0d44e2540f472c0635ab62c232b.cloudfront.net (CloudFront)"
              }
            ]
          }
        }
      }
    }
  ]
}
```

```

    }
  ],
  "host": [
    {
      "key": "Host",
      "value": "example.org"
    }
  ],
  "cache-control": [
    {
      "key": "Cache-Control",
      "value": "no-cache, cf-no-cache"
    }
  ]
},
"method": "GET",
"origin": {
  "custom": {
    "customHeaders": {},
    "domainName": "example.org",
    "keepaliveTimeout": 5,
    "path": "",
    "port": 443,
    "protocol": "https",
    "readTimeout": 30,
    "sslProtocols": [
      "TLSv1",
      "TLSv1.1",
      "TLSv1.2"
    ]
  }
},
"queryString": "",
"uri": "/"
}
}
}
]
}

```

Request Event Fields

Request event object data is contained in two subobjects: `config` (`Records.cf.config`) and `request` (`Records.cf.request`). The following lists describe each subobject's fields.

Fields in the Config Object

The following list describes the fields in the `config` object (`Records.cf.config`).

distributionDomainName (read-only)

The domain name of the distribution that's associated with the request.

distributionID (read-only)

The ID of the distribution that's associated with the request.

eventType (read-only)

The type of trigger that's associated with the request: `viewer-request` or `origin-request`.

requestId (read-only)

An encrypted string that uniquely identifies a viewer-to-CloudFront request. The `requestId` value also appears in CloudFront access logs as `x-edge-request-id`. For more information, see

[Configuring and Using Standard Logs \(Access Logs\) \(p. 439\)](#) and [Web Distribution Standard Log File Format \(p. 445\)](#).

Fields in the Request Object

The following list describes the fields in the `request` object (`Records.cf.request`).

clientIp (read-only)

The IP address of the viewer that made the request. If the viewer used an HTTP proxy or a load balancer to send the request, the value is the IP address of the proxy or load balancer.

headers (read/write)

The headers in the request. Note the following:

- The keys in the `headers` object are lowercase versions of standard HTTP header names. Using lowercase keys gives you case-insensitive access to the header values.
- Each header object (for example, `headers["accept"]` or `headers["host"]`) is an array of key-value pairs. For a given header, the array contains one key-value pair for each value in the request.
- `key` contains the case-sensitive name of the header as it appeared in the HTTP request; for example, `Host`, `User-Agent`, `X-Forwarded-For`, and so on.
- `value` contains the header value as it appeared in the HTTP request.
- When your Lambda function adds or modifies request headers and you don't include the header `key` field, Lambda@Edge automatically inserts a header `key` using the header name that you provide. Regardless of how you've formatted the header name, the header `key` that's inserted automatically is formatted with initial capitalization for each part, separated by hyphens (-).

For example, you can add a header like the following, without a header `key`:

```
"user-agent": [  
  {  
    "value": "ExampleCustomUserAgent/1.X.0"  
  }  
]
```

In this example, Lambda@Edge automatically inserts `"key": "User-Agent"`.

For information about restrictions on header usage, see [Headers \(p. 394\)](#).

method (read-only)

The HTTP method of the request.

queryString (read/write)

The query string, if any, in the request. If the request doesn't include a query string, the event object still includes `queryString` with an empty value. For more information about query strings, see [Caching Content Based on Query String Parameters \(p. 241\)](#).

uri (read/write)

The relative path of the requested object. If your Lambda function modifies the `uri` value, note the following:

- The new `uri` value must begin with a forward slash (/).
- When a function changes the `uri` value, that changes the object that the viewer is requesting.
- When a function changes the `uri` value, that *doesn't* change the cache behavior for the request or the origin that the request is sent to.

body (read/write)

The body of the HTTP request. The body structure can contain the following fields:

inputTruncated (read-only)

A Boolean flag that indicates whether the body was truncated by Lambda@Edge. For more information, see [Size Quotas on Request Body with the Include Body Option \(p. 398\)](#).

action (read/write)

The action that you intend to take with the body. The options for `action` are the following:

- `read-only`: This is the default. When returning the response from the Lambda function, if `action` is `read-only`, Lambda@Edge ignores any changes to `encoding` or `data`.
- `replace`: Specify this when you want to replace the body sent to the origin.

encoding (read/write)

The encoding for the body. When Lambda@Edge exposes the body to the Lambda function, it first converts the body to base64-encoding. If you choose `replace` for the `action` to replace the body, you can opt to use `base64` (the default) or `text` encoding. If you specify `encoding` as `base64` but the body is not valid base64, CloudFront returns an error.

data (read/write)

The request body content.

origin (read/write) (origin events only)

The origin to send the request to. The `origin` structure must contain exactly one origin, which can be a custom origin or an Amazon S3 origin. The origin structure can contain the following fields:

customHeaders (read/write) (custom and Amazon S3 origins)

You can include custom headers with the request by specifying a header name and value pair for each custom header. You can't add headers that are blacklisted, and a header with the same name can't be present in `Records.cf.request.headers`. The [notes about request headers \(p. 356\)](#) also apply to custom headers. For more information, see [Custom Headers that CloudFront Can't Add to Origin Requests \(p. 284\)](#) and [Blacklisted Headers \(p. 394\)](#).

domainName (read/write) (custom and Amazon S3 origins)

The domain name of the origin. The domain name can't be empty.

- **For custom origins** – Specify a DNS domain name, such as `www.example.com`. The domain name can't include a colon (:), and can't be an IP address. The domain name can be up to 253 characters.
- **For Amazon S3 origins** – Specify the DNS domain name of the Amazon S3 bucket, such as `awsexamplebucket.s3.eu-west-1.amazonaws.com`. The name can be up to 128 characters, and must be all lowercase.

path (read/write) (custom and Amazon S3 origins)

The directory path at the origin where the request should locate content. The path should start with a forward slash (/) but shouldn't end with one (for example, it shouldn't end with `example-path/`). For custom origins only, the path should be URL encoded and have a maximum length of 255 characters.

keepaliveTimeout (read/write) (custom origins only)

How long, in seconds, that CloudFront should try to maintain the connection to the origin after receiving the last packet of the response. The value must be a number from 1–60, inclusive.

port (read/write) (custom origins only)

The port that CloudFront should connect to at your custom origin. The port must be 80, 443, or a number in the range of 1024–65535, inclusive.

protocol (read/write) (custom origins only)

The connection protocol that CloudFront should use when connecting to your origin. The value can be `http` or `https`.

readTimeout (read/write) (custom origins only)

How long, in seconds, CloudFront should wait for a response after sending a request to your origin. This also specifies how long CloudFront should wait after receiving a packet of a response before receiving the next packet. The value must be a number from 4–60, inclusive.

sslProtocols (read/write) (custom origins only)

The minimum SSL/TLS protocol that CloudFront can use when establishing an HTTPS connection with your origin. Values can be any of the following: `TLSv1.2`, `TLSv1.1`, `TLSv1`, or `SSLv3`.

authMethod (read/write) (Amazon S3 origins only)

If you're using an [origin access identity \(OAI\)](#) (p. 209), set this field to `origin-access-identity`, or set it to `none` if you aren't using an OAI. If you set `authMethod` to `origin-access-identity`, there are several requirements:

- You must specify the `region` (see the following field).
- You must use the same OAI when you change the request from one Amazon S3 origin to another.
- You can't use an OAI when you change the request from a custom origin to an Amazon S3 origin.

region (read/write) (Amazon S3 origins only)

The AWS Region of your Amazon S3 bucket. This is required only when you set `authMethod` to `origin-access-identity`.

Response Events

The following topics show the structure of the object that CloudFront passes to a Lambda function for [viewer and origin response events](#) (p. 339). Following the examples is a list of all the possible fields in viewer and origin response events.

Topics

- [Example Origin Response](#) (p. 358)
- [Example Viewer Response](#) (p. 360)
- [Response Event Fields](#) (p. 362)

Example Origin Response

The following example shows an origin response event object.

```
{
  "Records": [
    {
      "cf": {
        "config": {
          "distributionDomainName": "d111111abcdef8.cloudfront.net",
          "distributionId": "EDFDVBD6EXAMPLE",
          "eventType": "origin-response",
          "requestId": "4TyzHTaYwb1GX1qTfsHhEqV6HUdd_BzoBZnwfnvQc_1oF26ClkoUSEQ=="
        },

```

```
"request": {
  "clientIp": "203.0.113.178",
  "headers": {
    "x-forwarded-for": [
      {
        "key": "X-Forwarded-For",
        "value": "203.0.113.178"
      }
    ],
    "user-agent": [
      {
        "key": "User-Agent",
        "value": "Amazon CloudFront"
      }
    ],
    "via": [
      {
        "key": "Via",
        "value": "2.0 8f22423015641505b8c857a37450d6c0.cloudfront.net (CloudFront)"
      }
    ],
    "host": [
      {
        "key": "Host",
        "value": "example.org"
      }
    ],
    "cache-control": [
      {
        "key": "Cache-Control",
        "value": "no-cache, cf-no-cache"
      }
    ]
  },
  "method": "GET",
  "origin": {
    "custom": {
      "customHeaders": {},
      "domainName": "example.org",
      "keepaliveTimeout": 5,
      "path": "",
      "port": 443,
      "protocol": "https",
      "readTimeout": 30,
      "sslProtocols": [
        "TLSv1",
        "TLSv1.1",
        "TLSv1.2"
      ]
    }
  },
  "queryString": "",
  "uri": "/"
},
"response": {
  "headers": {
    "access-control-allow-credentials": [
      {
        "key": "Access-Control-Allow-Credentials",
        "value": "true"
      }
    ],
    "access-control-allow-origin": [
      {
        "key": "Access-Control-Allow-Origin",
        "value": "*"
      }
    ]
  }
}
```



```
"Records": [
  {
    "cf": {
      "config": {
        "distributionDomainName": "d111111abcdef8.cloudfront.net",
        "distributionId": "EDFDVBD6EXAMPLE",
        "eventType": "viewer-response",
        "requestId": "4TyzHTaYWb1GX1qTfsHhEqV6HUDd_BzoBZnwfnvQc_1oF26ClkoUSEQ=="
      },
      "request": {
        "clientIp": "203.0.113.178",
        "headers": {
          "host": [
            {
              "key": "Host",
              "value": "d111111abcdef8.cloudfront.net"
            }
          ],
          "user-agent": [
            {
              "key": "User-Agent",
              "value": "curl/7.66.0"
            }
          ],
          "accept": [
            {
              "key": "accept",
              "value": "*/*"
            }
          ]
        },
        "method": "GET",
        "queryString": "",
        "uri": "/"
      },
      "response": {
        "headers": {
          "access-control-allow-credentials": [
            {
              "key": "Access-Control-Allow-Credentials",
              "value": "true"
            }
          ],
          "access-control-allow-origin": [
            {
              "key": "Access-Control-Allow-Origin",
              "value": "*"
            }
          ],
          "date": [
            {
              "key": "Date",
              "value": "Mon, 13 Jan 2020 20:14:56 GMT"
            }
          ],
          "referrer-policy": [
            {
              "key": "Referrer-Policy",
              "value": "no-referrer-when-downgrade"
            }
          ],
          "server": [
            {
              "key": "Server",
              "value": "ExampleCustomOriginServer"
            }
          ]
        }
      }
    }
  }
]
```

```

    ],
    "x-content-type-options": [
      {
        "key": "X-Content-Type-Options",
        "value": "nosniff"
      }
    ],
    "x-frame-options": [
      {
        "key": "X-Frame-Options",
        "value": "DENY"
      }
    ],
    "x-xss-protection": [
      {
        "key": "X-XSS-Protection",
        "value": "1; mode=block"
      }
    ],
    "age": [
      {
        "key": "Age",
        "value": "2402"
      }
    ],
    "content-type": [
      {
        "key": "Content-Type",
        "value": "text/html; charset=utf-8"
      }
    ],
    "content-length": [
      {
        "key": "Content-Length",
        "value": "9593"
      }
    ]
  },
  "status": "200",
  "statusDescription": "OK"
}
}
}
]
}

```

Response Event Fields

Response event object data is contained in three subobjects: `config` (`Records.cf.config`), `request` (`Records.cf.request`), and `response` (`Records.cf.response`). For more information about the fields in the request object, see [Fields in the Request Object \(p. 356\)](#). The following lists describe the fields in the `config` and `response` subobjects.

Fields in the Config Object

The following list describes the fields in the `config` object (`Records.cf.config`).

distributionDomainName (read-only)

The domain name of the distribution that's associated with the response.

distributionID (read-only)

The ID of the distribution that's associated with the response.

eventType (read-only)

The type of trigger that's associated with the response: `origin-response` or `viewer-response`.

requestId (read-only)

An encrypted string that uniquely identifies the viewer-to-CloudFront request that this response is associated with. The `requestId` value also appears in CloudFront access logs as `x-edge-request-id`. For more information, see [Configuring and Using Standard Logs \(Access Logs\) \(p. 439\)](#) and [Web Distribution Standard Log File Format \(p. 445\)](#).

Fields in the Response Object

The following list describes the fields in the `response` object (`Records.cf.response`). For information about using a `Lambda@Edge` function to generate an HTTP response, see [Generating HTTP Responses in Request Triggers \(p. 364\)](#).

headers (read/write)

The headers in the response. Note the following:

- The keys in the `headers` object are lowercase versions of standard HTTP header names. Using lowercase keys gives you case-insensitive access to the header values.
- Each header object (for example, `headers["content-type"]` or `headers["content-length"]`) is an array of key-value pairs. For a given header, the array contains one key-value pair for each value in the response.
- `key` contains the case-sensitive name of the header as it appears in the HTTP response; for example, `Content-Type`, `Content-Length`, and so on.
- `value` contains the header value as it appears in the HTTP response.
- When your Lambda function adds or modifies response headers and you don't include the header `key` field, `Lambda@Edge` automatically inserts a header key using the header name that you provide. Regardless of how you've formatted the header name, the header key that's inserted automatically is formatted with initial capitalization for each part, separated by hyphens (-).

For example, you can add a header like the following, without a header key:

```
"content-type": [  
  {  
    "value": "text/html;charset=UTF-8"  
  }  
]
```

In this example, `Lambda@Edge` automatically inserts `"key": "Content-Type"`.

For information about restrictions on header usage, see [Headers \(p. 394\)](#).

status

The HTTP status code of the response.

statusDescription

The HTTP status description of the response.

Working with Requests and Responses

You can use `Lambda@Edge` to work with requests and responses in several ways:

- [Generating HTTP Responses in Request Triggers](#) (p. 364)
- [Updating HTTP Responses in Origin-Response Triggers](#) (p. 366)
- [Accessing the Request Body by Choosing the Include Body Option](#) (p. 367)
- [Using Lambda@Edge Functions with Origin Failover](#) (p. 364)

Using Lambda@Edge Functions with Origin Failover

You can use Lambda@Edge functions with CloudFront distributions that you've set up with origin groups, for example, for origin failover that you configure to help ensure high availability. To use a Lambda function with an origin group, specify the function in an origin request or origin response trigger for an origin group when you create the cache behavior.

For more information, see the following:

- **Creating origin groups:** [Creating an Origin Group](#) (p. 232)
- **How origin failover works with Lambda@Edge:** [Use Origin Failover with Lambda@Edge Functions](#) (p. 234)

Generating HTTP Responses in Request Triggers

When CloudFront receives a request, you can use a Lambda function to generate an HTTP response that CloudFront returns directly to the viewer without forwarding the response to the origin. Generating HTTP responses reduces the load on the origin, and typically also reduces latency for the viewer.

Some common scenarios for generating HTTP responses include the following:

- Returning a small webpage to the viewer
- Returning an HTTP 301 or 302 status code to redirect the user to another webpage
- Returning an HTTP 401 status code to the viewer when the user hasn't authenticated

A Lambda@Edge function can generate an HTTP response when the following CloudFront events occur:

Viewer request events

When a function is triggered by a viewer request event, CloudFront returns the response to the viewer and doesn't cache it.

Origin request events

When a function is triggered by an origin request event, CloudFront checks the edge cache for a response that was previously generated by the function.

- If the response is in the cache, the function isn't executed and CloudFront returns the cached response to the viewer.
- If the response isn't in the cache, the function is executed, CloudFront returns the response to the viewer, and also caches it.

To see some sample code for generating HTTP responses, see [Lambda@Edge Example Functions](#) (p. 367). You can also replace the HTTP responses in response triggers. For more information, see [Updating HTTP Responses in Origin-Response Triggers](#) (p. 366).

Programming Model

This section describes the programming model for using Lambda@Edge to generate HTTP responses.

Topics

- [Response Object \(p. 365\)](#)
- [Errors \(p. 366\)](#)
- [Required Fields \(p. 366\)](#)

Response Object

The response you return as the `result` parameter of the `callback` method should have the following structure (note that only the `status` field is required).

```
const response = {
  body: 'content',
  bodyEncoding: 'text' | 'base64',
  headers: {
    'header name in lowercase': [{
      key: 'header name in standard case',
      value: 'header value'
    }],
    ...
  },
  status: 'HTTP status code',
  statusDescription: 'status description'
};
```

The response object can include the following values:

body

The body, if any, that you want CloudFront to return in the generated response.

bodyEncoding

The encoding for the value that you specified in the `body`. The only valid encodings are `text` and `base64`. If you include `body` in the response object but omit `bodyEncoding`, CloudFront treats the body as text.

If you specify `bodyEncoding` as `base64` but the body is not valid base64, CloudFront returns an error.

headers

Headers that you want CloudFront to return in the generated response. Note the following:

- The keys in the `headers` object are lowercase versions of standard HTTP header names. Using lowercase keys gives you case-insensitive access to the header values.
- Each header (for example, `headers["accept"]` or `headers["host"]`) is an array of key-value pairs. For a given header, the array contains one key-value pair for each value in the generated response.
- `key` (optional) is the case-sensitive name of the header as it appears in an HTTP request; for example, `accept` or `host`.
- Specify `value` as a header value.
- If you do not include the header key portion of the key-value pair, Lambda@Edge automatically inserts a header key using the header name that you provide. Regardless of how you've formatted the header name, the header key that is inserted is automatically formatted with initial capitalization for each part, separated by hyphens (-).

For example, you can add a header like the following, without a header key: `'content-type'` :
`[{ value: 'text/html;charset=UTF-8' }]`

In this example, Lambda@Edge creates the following header key: `Content-Type`.

For information about restrictions on header usage, see [Headers \(p. 394\)](#).

status

The HTTP status code that you want CloudFront to use for the following:

- Return in the response
- Cache in the CloudFront edge cache, when the response was generated by a function that was triggered by an origin request event
- Log in CloudFront [Configuring and Using Standard Logs \(Access Logs\) \(p. 439\)](#)

If the `status` value isn't between 200 and 599, CloudFront returns an error to the viewer.

statusDescription

The description that you want CloudFront to return in the response, to accompany the HTTP status code. You don't need to use standard descriptions, such as `OK` for an HTTP status code of 200.

Errors

The following are possible errors for generated HTTP responses.

Response Contains a Body and Specifies 204 (No Content) for Status

When a function is triggered by a viewer request, CloudFront returns an HTTP 502 status code (Bad Gateway) to the viewer when both of the following are true:

- The value of `status` is 204 (No Content)
- The response includes a value for `body`

This is because Lambda@Edge imposes the optional restriction found in RFC 2616, which states that an `HTTP 204` response does not need to contain a message body.

Restrictions on the Size of the Generated Response

The maximum size of a response that is generated by a Lambda function depends on the event that triggered the function:

- **Viewer request events** – 40 KB
- **Origin request events** – 1 MB

If the response is larger than the allowed size, CloudFront returns an HTTP 502 status code (Bad Gateway) to the viewer.

Required Fields

The `status` field is required.

All other fields are optional.

Updating HTTP Responses in Origin-Response Triggers

When CloudFront receives an HTTP response from the origin server, if there is an origin-response trigger associated with the cache behavior, you can modify the HTTP response to override what was returned from the origin.

Some common scenarios for updating HTTP responses include the following:

- Changing the status to set an HTTP 200 status code and creating static body content to return to the viewer when an origin returns an error status code (4xx or 5xx). For sample code, see [Example: Using an Origin-Response Trigger to Update the Error Status Code to 200-OK](#) (p. 388).
- Changing the status to set an HTTP 301 or HTTP 302 status code, to redirect the user to another website when an origin returns an error status code (4xx or 5xx). For sample code, see [Example: Using an Origin-Response Trigger to Update the Error Status Code to 302-Found](#) (p. 389).

You can also replace the HTTP responses in viewer and origin request events. For more information, see [Generating HTTP Responses in Request Triggers](#) (p. 364).

When you're working with the HTTP response, Lambda@Edge does not expose the HTML body that is returned by the origin server to the origin-response trigger. You can generate a static content body by setting it to the desired value, or remove the body inside the function by setting the value to be empty. If you don't update the body field in your function, the original body returned by the origin server is returned back to viewer.

Accessing the Request Body by Choosing the Include Body Option

You can opt to have Lambda@Edge expose the body in a request for writable HTTP methods (POST, PUT, DELETE, and so on), so that you can access it in your Lambda function. You can choose read-only access, or you can specify that you'll replace the body.

To enable this option, choose **Include Body** when you create a CloudFront trigger for your function that's for a viewer request or origin request event. For more information, see [Adding Triggers for a Lambda@Edge Function](#) (p. 339), or to learn about using **Include Body** with your function, see [Lambda@Edge Event Structure](#) (p. 353).

Scenarios when you might want to use this feature include the following:

- Processing web forms, like "contact us" forms, without sending customer input data back to origin servers.
- Gathering web beacon data that's sent by viewer browsers and processing it at the edge.

For sample code, see [Lambda@Edge Example Functions](#) (p. 367).

Note

If the request body is large, Lambda@Edge truncates it. For detailed information about the maximum size and truncation, see [Maximum Size for Body with the Include Body Option](#) (p. 399).

Lambda@Edge Example Functions

See the following sections for examples of using Lambda functions with CloudFront.

Note

For Node.js functions, each function must call the `callback` parameter to successfully process a request or return a response. For more information, see [Writing and Creating a Lambda@Edge Function](#) (p. 335).

Topics

- [General Examples \(p. 368\)](#)
- [Generating Responses - Examples \(p. 371\)](#)
- [Working with Query Strings - Examples \(p. 373\)](#)
- [Personalize Content by Country or Device Type Headers - Examples \(p. 377\)](#)
- [Content-Based Dynamic Origin Selection - Examples \(p. 381\)](#)
- [Updating Error Statuses - Examples \(p. 388\)](#)
- [Accessing the Request Body - Examples \(p. 390\)](#)

General Examples

The examples in this section illustrate some common ways to use Lambda@Edge in CloudFront.

Topics

- [Example: A/B Testing \(p. 368\)](#)
- [Example: Overriding a Response Header \(p. 370\)](#)

Example: A/B Testing

You can use the following example to test two different versions of an image without creating redirects or changing the URL. This example reads the cookies in the viewer request and modifies the request URL accordingly. If the viewer doesn't send a cookie with one of the expected values, the example randomly assigns the viewer to one of the URLs.

Node.js

```
'use strict';

exports.handler = (event, context, callback) => {
  const request = event.Records[0].cf.request;
  const headers = request.headers;

  if (request.uri !== '/experiment-pixel.jpg') {
    // do not process if this is not an A-B test request
    callback(null, request);
    return;
  }

  const cookieExperimentA = 'X-Experiment-Name=A';
  const cookieExperimentB = 'X-Experiment-Name=B';
  const pathExperimentA = '/experiment-group/control-pixel.jpg';
  const pathExperimentB = '/experiment-group/treatment-pixel.jpg';

  /*
   * Lambda at the Edge headers are array objects.
   *
   * Client may send multiple Cookie headers, i.e.:
   * > GET /viewerRes/test HTTP/1.1
   * > User-Agent: curl/7.18.1 (x86_64-unknown-linux-gnu) libcurl/7.18.1
   * > Cookie: First=1; Second=2
   * > Cookie: ClientCode=abc
   * > Host: example.com
   *
   * You can access the first Cookie header at headers["cookie"][0].value
   * and the second at headers["cookie"][1].value.
   *
   * Header values are not parsed. In the example above,
```

```
* headers["cookie"][0].value is equal to "First=1; Second=2"
*/
let experimentUri;
if (headers.cookie) {
  for (let i = 0; i < headers.cookie.length; i++) {
    if (headers.cookie[i].value.indexOf(cookieExperimentA) >= 0) {
      console.log('Experiment A cookie found');
      experimentUri = pathExperimentA;
      break;
    } else if (headers.cookie[i].value.indexOf(cookieExperimentB) >= 0) {
      console.log('Experiment B cookie found');
      experimentUri = pathExperimentB;
      break;
    }
  }
}

if (!experimentUri) {
  console.log('Experiment cookie has not been found. Throwing dice...');
  if (Math.random() < 0.75) {
    experimentUri = pathExperimentA;
  } else {
    experimentUri = pathExperimentB;
  }
}

request.uri = experimentUri;
console.log(`Request uri set to "${request.uri}"`);
callback(null, request);
};
```

Python

```
import json
import random

def lambda_handler(event, context):
    request = event['Records'][0]['cf']['request']
    headers = request['headers']

    if request['uri'] != '/experiment-pixel.jpg':
        # Not an A/B Test
        return request

    cookieExperimentA, cookieExperimentB = 'X-Experiment-Name=A', 'X-Experiment-Name=B'
    pathExperimentA, pathExperimentB = '/experiment-group/control-pixel.jpg', '/experiment-group/treatment-pixel.jpg'

    '''
    Lambda at the Edge headers are array objects.

    Client may send multiple cookie headers. For example:
    > GET /viewerRes/test HTTP/1.1
    > User-Agent: curl/7.18.1 (x86_64-unknown-linux-gnu) libcurl/7.18.1 OpenSSL/1.0.1u
    zlib/1.2.3
    > Cookie: First=1; Second=2
    > Cookie: ClientCode=abc
    > Host: example.com

    You can access the first Cookie header at headers["cookie"][0].value
    and the second at headers["cookie"][1].value.

    Header values are not parsed. In the example above,
    headers["cookie"][0].value is equal to "First=1; Second=2"
    '''
```

```
'''
experimentUri = ""

for cookie in headers.get('cookie', []):
    if cookieExperimentA in cookie['value']:
        print("Experiment A cookie found")
        experimentUri = pathExperimentA
        break
    elif cookieExperimentB in cookie['value']:
        print("Experiment B cookie found")
        experimentUri = pathExperimentB
        break

if not experimentUri:
    print("Experiment cookie has not been found. Throwing dice...")
    if random.random() < 0.75:
        experimentUri = pathExperimentA
    else:
        experimentUri = pathExperimentB

request['uri'] = experimentUri
print(f"Request uri set to {experimentUri}")
return request
```

Example: Overriding a Response Header

The following example shows how to change the value of a response header based on the value of another header.

Node.js

```
'use strict';

exports.handler = (event, context, callback) => {
    const response = event.Records[0].cf.response;
    const headers = response.headers;

    const headerNameSrc = 'X-Amz-Meta-Last-Modified';
    const headerNameDst = 'Last-Modified';

    if (headers[headerNameSrc.toLowerCase()]) {
        headers[headerNameDst.toLowerCase()] = [
            headers[headerNameSrc.toLowerCase()][0],
        ];
        console.log(`Response header "${headerNameDst}" was set to ` +
            `"${headers[headerNameDst.toLowerCase()][0].value}"`);
    }

    callback(null, response);
};
```

Python

```
import json

def lambda_handler(event, context):
    response = event["Records"][0]["cf"]["response"]
    headers = response["headers"]

    headerNameSrc = "X-Amz-Meta-Last-Modified"
```

```
headerNameDst = "Last-Modified"

if headers.get(headerNameSrc.lower(), None):
    headers[headerNameDst.lower()] = [headers[headerNameSrc.lower()][0]]
    print(f"Response header {headerNameDst.lower()} was set to {headers[headerNameSrc.lower()][0]}")

return response
```

Generating Responses - Examples

The examples in this section show how you can use Lambda@Edge to generate responses.

Topics

- [Example: Serving Static Content \(Generated Response\) \(p. 371\)](#)
- [Example: Generating an HTTP Redirect \(Generated Response\) \(p. 372\)](#)

Example: Serving Static Content (Generated Response)

The following example shows how to use a Lambda function to serve static website content, which reduces the load on the origin server and reduces overall latency.

Note

You can generate HTTP responses for viewer request and origin request events. For more information, see [Generating HTTP Responses in Request Triggers \(p. 364\)](#). You can also replace the HTTP response in origin and viewer response events. For more information, see [Updating HTTP Responses in Origin-Response Triggers \(p. 366\)](#).

Node.js

```
'use strict';

const content = `
<!DOCTYPE html>
<html lang="en">
  <head>
    <meta charset="utf-8">
    <title>Simple Lambda@Edge Static Content Response</title>
  </head>
  <body>
    <p>Hello from Lambda@Edge!</p>
  </body>
</html>
`;

exports.handler = (event, context, callback) => {
  /*
   * Generate HTTP OK response using 200 status code with HTML body.
   */
  const response = {
    status: '200',
    statusDescription: 'OK',
    headers: {
      'cache-control': [{
        key: 'Cache-Control',
        value: 'max-age=100'
      }],
      'content-type': [{
        key: 'Content-Type',
```

```
        value: 'text/html'
    }],
    body: content,
};
callback(null, response);
};
```

Python

```
import json

CONTENT = """
<!DOCTYPE html>
<html lang="en">
  <head>
    <meta charset="utf-8">
    <title>Simple Lambda@Edge Static Content Response</title>
  </head>
  <body>
    <p>Hello from Lambda@Edge!</p>
  </body>
</html>
"""

def lambda_handler(event, context):
    # Generate HTTP OK response using 200 status code with HTML body.
    response = {
        'status': '200',
        'statusDescription': 'OK',
        'headers': {
            'cache-control': [
                {
                    'key': 'Cache-Control',
                    'value': 'max-age=100'
                }
            ],
            "content-type": [
                {
                    'key': 'Content-Type',
                    'value': 'text/html'
                }
            ]
        },
        'body': CONTENT
    }
    return response
```

Example: Generating an HTTP Redirect (Generated Response)

The following example shows how to generate an HTTP redirect.

Note

You can generate HTTP responses for viewer request and origin request events. For more information, see [Generating HTTP Responses in Request Triggers \(p. 364\)](#).

Node.js

```
'use strict';

exports.handler = (event, context, callback) => {
    /*
```

```
* Generate HTTP redirect response with 302 status code and Location header.
*/
const response = {
  status: '302',
  statusDescription: 'Found',
  headers: {
    location: [{
      key: 'Location',
      value: 'http://docs.aws.amazon.com/lambda/latest/dg/lambda-edge.html',
    }],
  },
};
callback(null, response);
};
```

Python

```
def lambda_handler(event, context):

    # Generate HTTP redirect response with 302 status code and Location header.

    response = {
        'status': '302',
        'statusDescription': 'Found',
        'headers': {
            'location': [{
                'key': 'Location',
                'value': 'http://docs.aws.amazon.com/lambda/latest/dg/lambda-
edge.html'
            }]
        }
    }

    return response
```

Working with Query Strings - Examples

The examples in this section include ways that you can use Lambda@Edge with query strings.

Topics

- [Example: Adding a Header Based on a Query String Parameter \(p. 373\)](#)
- [Example: Normalizing Query String Parameters to Improve the Cache Hit Ratio \(p. 374\)](#)
- [Example: Redirecting Unauthenticated Users to a Sign-In Page \(p. 376\)](#)

Example: Adding a Header Based on a Query String Parameter

The following example shows how to get the key-value pair of a query string parameter, and then add a header based on those values.

Node.js

```
'use strict';

const querystring = require('querystring');
exports.handler = (event, context, callback) => {
  const request = event.Records[0].cf.request;

  /* When a request contains a query string key-value pair but the origin server
```

```
    * expects the value in a header, you can use this Lambda function to
    * convert the key-value pair to a header. Here's what the function does:
    * 1. Parses the query string and gets the key-value pair.
    * 2. Adds a header to the request using the key-value pair that the function got
in step 1.
    */

    /* Parse request querystring to get javascript object */
    const params = querystring.parse(request.querystring);

    /* Move auth param from querystring to headers */
    const headerName = 'Auth-Header';
    request.headers[headerName.toLowerCase()] = [{ key: headerName, value:
params.auth }];
    delete params.auth;

    /* Update request querystring */
    request.querystring = querystring.stringify(params);

    callback(null, request);
}

;
```

Python

```
from urllib.parse import parse_qs, urlencode

def lambda_handler(event, context):
    request = event['Records'][0]['cf']['request']

    '''
    When a request contains a query string key-value pair but the origin server
    expects the value in a header, you can use this Lambda function to
    convert the key-value pair to a header. Here's what the function does:
        1. Parses the query string and gets the key-value pair.
        2. Adds a header to the request using the key-value pair that the function got
in step 1.
    '''

    # Parse request querystring to get dictionary/json
    params = {k : v[0] for k, v in parse_qs(request['querystring']).items()}

    # Move auth param from querystring to headers
    headerName = 'Auth-Header'
    request['headers'][headerName.lower()] = [{'key': headerName, 'value':
params['auth']}]}
    del params['auth']

    # Update request querystring
    request['querystring'] = urlencode(params)

    return request
```

Example: Normalizing Query String Parameters to Improve the Cache Hit Ratio

The following example shows how to improve your cache hit ratio by making the following changes to query strings before CloudFront forwards requests to your origin:

- Alphabetize key-value pairs by the name of the parameter.

- Change the case of key-value pairs to lowercase.

For more information, see [Caching Content Based on Query String Parameters \(p. 241\)](#).

Node.js

```
'use strict';

const querystring = require('querystring');

exports.handler = (event, context, callback) => {
  const request = event.Records[0].cf.request;
  /* When you configure a distribution to forward query strings to the origin and
   * to cache based on a whitelist of query string parameters, we recommend
   * the following to improve the cache-hit ratio:
   * - Always list parameters in the same order.
   * - Use the same case for parameter names and values.
   *
   * This function normalizes query strings so that parameter names and values
   * are lowercase and parameter names are in alphabetical order.
   *
   * For more information, see:
   * http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/
   QueryStringParameters.html
   */

  console.log('Query String: ', request.querystring);

  /* Parse request query string to get javascript object */
  const params = querystring.parse(request.querystring.toLowerCase());
  const sortedParams = {};

  /* Sort param keys */
  Object.keys(params).sort().forEach(key => {
    sortedParams[key] = params[key];
  });

  /* Update request querystring with normalized */
  request.querystring = querystring.stringify(sortedParams);

  callback(null, request);
};
```

Python

```
from urllib.parse import parse_qs, urlencode

def lambda_handler(event, context):
    request = event['Records'][0]['cf']['request']
    '''
    When you configure a distribution to forward query strings to the origin and
    to cache based on a whitelist of query string parameters, we recommend
    the following to improve the cache-hit ratio:
    Always list parameters in the same order.
    - Use the same case for parameter names and values.

    This function normalizes query strings so that parameter names and values
    are lowercase and parameter names are in alphabetical order.

    For more information, see:
    http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/
    QueryStringParameters.html
    '''
```



```
print("Query string: ", request["querystring"])

# Parse request query string to get js object
params = {k : v[0] for k, v in parse_qs(request['querystring'].lower()).items()}

# Sort param keys
sortedParams = sorted(params.items(), key=lambda x: x[0])

# Update request querystring with normalized
request['querystring'] = urlencode(sortedParams)

return request
```

Example: Redirecting Unauthenticated Users to a Sign-In Page

The following example shows how to redirect users to a sign-in page if they haven't entered their credentials.

Node.js

```
'use strict';

function parseCookies(headers) {
  const parsedCookie = {};
  if (headers.cookie) {
    headers.cookie[0].value.split(';').forEach((cookie) => {
      if (cookie) {
        const parts = cookie.split('=');
        parsedCookie[parts[0].trim()] = parts[1].trim();
      }
    });
  }
  return parsedCookie;
}

exports.handler = (event, context, callback) => {
  const request = event.Records[0].cf.request;
  const headers = request.headers;

  /* Check for session-id in request cookie in viewer-request event,
   * if session-id is absent, redirect the user to sign in page with original
   * request sent as redirect_url in query params.
   */

  /* Check for session-id in cookie, if present then proceed with request */
  const parsedCookies = parseCookies(headers);
  if (parsedCookies && parsedCookies['session-id']) {
    callback(null, request);
    return;
  }

  /* URI encode the original request to be sent as redirect_url in query params */
  const encodedRedirectUrl = encodeURIComponent(`https://${headers.host[0].value}${request.uri}?${request.querystring}`);
  const response = {
    status: '302',
    statusDescription: 'Found',
    headers: {
      location: [{
        key: 'Location',
        value: `https://www.example.com/signin?redirect_url=${encodedRedirectUrl}`,
      }],
    },
  };
}
```

```
    },  
    };  
    callback(null, response);  
};
```

Python

```
import urllib  
  
def parseCookies(headers):  
    parsedCookie = {}  
    if headers.get('cookie'):  
        for cookie in headers['cookie'][0]['value'].split(';'):  
            if cookie:  
                parts = cookie.split('=')  
                parsedCookie[parts[0].strip()] = parts[1].strip()  
    return parsedCookie  
  
def lambda_handler(event, context):  
    request = event['Records'][0]['cf']['request']  
    headers = request['headers']  
  
    '''  
    Check for session-id in request cookie in viewer-request event,  
    if session-id is absent, redirect the user to sign in page with original  
    request sent as redirect_url in query params.  
    '''  
  
    # Check for session-id in cookie, if present, then proceed with request  
    parsedCookies = parseCookies(headers)  
  
    if parsedCookies and parsedCookies['session-id']:  
        return request  
  
    # URI encode the original request to be sent as redirect_url in query params  
    redirectUrl = "https://%s%s?%s" % (headers['host'][0]['value'], request['uri'],  
    request['querystring'])  
    encodedRedirectUrl = urllib.parse.quote_plus(redirectUrl.encode('utf-8'))  
  
    response = {  
        'status': '302',  
        'statusDescription': 'Found',  
        'headers': {  
            'location': [{  
                'key': 'Location',  
                'value': 'https://www.example.com/signin?redirect_url=%s' %  
    encodedRedirectUrl  
            }]  
        }  
    }  
    return response
```

Personalize Content by Country or Device Type Headers - Examples

The examples in this section illustrate how you can use Lambda@Edge to customize behavior based on location or the type of device used by the viewer.

Topics

- [Example: Redirecting Viewer Requests to a Country-Specific URL \(p. 378\)](#)

- [Example: Serving Different Versions of an Object Based on the Device \(p. 379\)](#)

Example: Redirecting Viewer Requests to a Country-Specific URL

The following example shows how to generate an HTTP redirect response with a country-specific URL and return the response to the viewer. This is useful when you want to provide country-specific responses. For example:

- If you have country-specific subdomains, such as `us.example.com` and `tw.example.com`, you can generate a redirect response when a viewer requests `example.com`.
- If you're streaming video but you don't have rights to stream the content in a specific country, you can redirect users in that country to a page that explains why they can't view the video.

Note the following:

- You must configure your distribution to cache based on the `CloudFront-Viewer-Country` header. For more information, see [Cache Based on Selected Request Headers \(p. 50\)](#).
- CloudFront adds the `CloudFront-Viewer-Country` header after the viewer request event. To use this example, you must create a trigger for the origin request event.

Node.js

```
'use strict';

/* This is an origin request function */
exports.handler = (event, context, callback) => {
  const request = event.Records[0].cf.request;
  const headers = request.headers;

  /*
   * Based on the value of the CloudFront-Viewer-Country header, generate an
   * HTTP status code 302 (Redirect) response, and return a country-specific
   * URL in the Location header.
   * NOTE: 1. You must configure your distribution to cache based on the
   *        CloudFront-Viewer-Country header. For more information, see
   *        http://docs.aws.amazon.com/console/cloudfront/cache-on-selected-headers
   *        2. CloudFront adds the CloudFront-Viewer-Country header after the viewer
   *        request event. To use this example, you must create a trigger for the
   *        origin request event.
   */

  let url = 'https://example.com/';
  if (headers['cloudfront-viewer-country']) {
    const countryCode = headers['cloudfront-viewer-country'][0].value;
    if (countryCode === 'TW') {
      url = 'https://tw.example.com/';
    } else if (countryCode === 'US') {
      url = 'https://us.example.com/';
    }
  }

  const response = {
    status: '302',
    statusDescription: 'Found',
    headers: {
      location: [{
        key: 'Location',
        value: url,
      }],
    },
  },
```

```
    },  
    };  
    callback(null, response);  
};
```

Python

```
# This is an origin request function  
  
def lambda_handler(event, context):  
    request = event['Records'][0]['cf']['request']  
    headers = request['headers']  
  
    '''  
    Based on the value of the CloudFront-Viewer-Country header, generate an  
    HTTP status code 302 (Redirect) response, and return a country-specific  
    URL in the Location header.  
    NOTE: 1. You must configure your distribution to cache based on the  
           CloudFront-Viewer-Country header. For more information, see  
           http://docs.aws.amazon.com/console/cloudfront/cache-on-selected-headers  
           2. CloudFront adds the CloudFront-Viewer-Country header after the viewer  
           request event. To use this example, you must create a trigger for the  
           origin request event.  
    '''  
  
    url = 'https://example.com/'  
    viewerCountry = headers.get('cloudfront-viewer-country')  
    if viewerCountry:  
        countryCode = viewerCountry[0]['value']  
        if countryCode == 'TW':  
            url = 'https://tw.example.com/'  
        elif countryCode == 'US':  
            url = 'https://us.example.com/'  
  
    response = {  
        'status': '302',  
        'statusDescription': 'Found',  
        'headers': {  
            'location': [{  
                'key': 'Location',  
                'value': url  
            }]  
        }  
    }  
  
    return response
```

Example: Serving Different Versions of an Object Based on the Device

The following example shows how to serve different versions of an object based on the type of device that the user is using, for example, a mobile device or a tablet. Note the following:

- You must configure your distribution to cache based on the `CloudFront-Is-*-Viewer` headers. For more information, see [Cache Based on Selected Request Headers \(p. 50\)](#).
- CloudFront adds the `CloudFront-Is-*-Viewer` headers after the viewer request event. To use this example, you must create a trigger for the origin request event.

Node.js

```
'use strict';

/* This is an origin request function */
exports.handler = (event, context, callback) => {
  const request = event.Records[0].cf.request;
  const headers = request.headers;

  /*
   * Serve different versions of an object based on the device type.
   * NOTE: 1. You must configure your distribution to cache based on the
   *        CloudFront-Is-*Viewer headers. For more information, see
   *        the following documentation:
   *        http://docs.aws.amazon.com/console/cloudfront/cache-on-selected-headers
   *        http://docs.aws.amazon.com/console/cloudfront/cache-on-device-type
   *        2. CloudFront adds the CloudFront-Is-*Viewer headers after the viewer
   *        request event. To use this example, you must create a trigger for the
   *        origin request event.
   */

  const desktopPath = '/desktop';
  const mobilePath = '/mobile';
  const tabletPath = '/tablet';
  const smarttvPath = '/smarttv';

  if (headers['cloudfront-is-desktop-viewer']
    && headers['cloudfront-is-desktop-viewer'][0].value === 'true') {
    request.uri = desktopPath + request.uri;
  } else if (headers['cloudfront-is-mobile-viewer']
    && headers['cloudfront-is-mobile-viewer'][0].value === 'true') {
    request.uri = mobilePath + request.uri;
  } else if (headers['cloudfront-is-tablet-viewer']
    && headers['cloudfront-is-tablet-viewer'][0].value === 'true') {
    request.uri = tabletPath + request.uri;
  } else if (headers['cloudfront-is-smarttv-viewer']
    && headers['cloudfront-is-smarttv-viewer'][0].value === 'true') {
    request.uri = smarttvPath + request.uri;
  }
  console.log(`Request uri set to "${request.uri}"`);

  callback(null, request);
};
```

Python

```
# This is an origin request function
def lambda_handler(event, context):
    request = event['Records'][0]['cf']['request']
    headers = request['headers']

    '''
    Serve different versions of an object based on the device type.
    NOTE: 1. You must configure your distribution to cache based on the
           CloudFront-Is-*Viewer headers. For more information, see
           the following documentation:
           http://docs.aws.amazon.com/console/cloudfront/cache-on-selected-headers
           http://docs.aws.amazon.com/console/cloudfront/cache-on-device-type
           2. CloudFront adds the CloudFront-Is-*Viewer headers after the viewer
           request event. To use this example, you must create a trigger for the
           origin request event.
    '''

    desktopPath = '/desktop';
```

```
mobilePath = '/mobile';
tabletPath = '/tablet';
smarttvPath = '/smarttv';

if 'cloudfront-is-desktop-viewer' in headers and headers['cloudfront-is-desktop-viewer'][0]['value'] == 'true':
    request['uri'] = desktopPath + request['uri']
elif 'cloudfront-is-mobile-viewer' in headers and headers['cloudfront-is-mobile-viewer'][0]['value'] == 'true':
    request['uri'] = mobilePath + request['uri']
elif 'cloudfront-is-tablet-viewer' in headers and headers['cloudfront-is-tablet-viewer'][0]['value'] == 'true':
    request['uri'] = tabletPath + request['uri']
elif 'cloudfront-is-smarttv-viewer' in headers and headers['cloudfront-is-smarttv-viewer'][0]['value'] == 'true':
    request['uri'] = smarttvPath + request['uri']

print("Request uri set to %s" % request['uri'])

return request
```

Content-Based Dynamic Origin Selection - Examples

The examples in this section show how you can use Lambda@Edge to route to different origins based on information in the request.

Topics

- [Example: Using an Origin-Request Trigger to Change From a Custom Origin to an Amazon S3 Origin \(p. 381\)](#)
- [Example: Using an Origin-Request Trigger to Change the Amazon S3 Origin Region \(p. 382\)](#)
- [Example: Using an Origin-Request Trigger to Change From an Amazon S3 Origin to a Custom Origin \(p. 385\)](#)
- [Example: Using an Origin-Request Trigger to Gradually Transfer Traffic From One Amazon S3 Bucket to Another \(p. 386\)](#)
- [Example: Using an Origin-Request Trigger to Change the Origin Domain Name Based on the Country Header \(p. 387\)](#)

Example: Using an Origin-Request Trigger to Change From a Custom Origin to an Amazon S3 Origin

This function demonstrates how an origin-request trigger can be used to change from a custom origin to an Amazon S3 origin from which the content is fetched, based on request properties.

Node.js

```
'use strict';

const querystring = require('querystring');

exports.handler = (event, context, callback) => {
    const request = event.Records[0].cf.request;

    /**
     * Reads query string to check if S3 origin should be used, and
     * if true, sets S3 origin properties.
     */
```

```
const params = querystring.parse(request.querystring);

if (params['useS3Origin']) {
  if (params['useS3Origin'] === 'true') {
    const s3DomainName = 'my-bucket.s3.amazonaws.com';

    /* Set S3 origin fields */
    request.origin = {
      s3: {
        domainName: s3DomainName,
        region: '',
        authMethod: 'none',
        path: '',
        customHeaders: {}
      }
    };
    request.headers['host'] = [{ key: 'host', value: s3DomainName}];
  }
}

callback(null, request);
};
```

Python

```
from urllib.parse import parse_qs

def lambda_handler(event, context):
    request = event['Records'][0]['cf']['request']
    '''
    Reads query string to check if S3 origin should be used, and
    if true, sets S3 origin properties
    '''
    params = {k: v[0] for k, v in parse_qs(request['querystring']).items()}
    if params.get('useS3Origin') == 'true':
        s3DomainName = 'my-bucket.s3.amazonaws.com'

        # Set S3 origin fields
        request['origin'] = {
            's3': {
                'domainName': s3DomainName,
                'region': '',
                'authMethod': 'none',
                'path': '',
                'customHeaders': {}
            }
        }
        request['headers']['host'] = [{'key': 'host', 'value': s3DomainName}]
    return request
```

Example: Using an Origin-Request Trigger to Change the Amazon S3 Origin Region

This function demonstrates how an origin-request trigger can be used to change the Amazon S3 origin from which the content is fetched, based on request properties.

In this example, we use the value of the CloudFront-Viewer-Country header to update the S3 bucket domain name to a bucket in a Region that is closer to the viewer. This can be useful in several ways:

- It reduces latencies when the Region specified is nearer to the viewer's country.

- It provides data sovereignty by making sure that data is served from an origin that's in the same country that the request came from.

To use this example, you must do the following:

- Configure your distribution to cache based on the CloudFront-Viewer-Country header. For more information, see [Cache Based on Selected Request Headers \(p. 50\)](#).
- Create a trigger for this function in the origin request event. CloudFront adds the CloudFront-Viewer-Country header after the viewer request event, so to use this example, you must make sure that the function executes for an origin request.

Node.js

```
'use strict';

exports.handler = (event, context, callback) => {
  const request = event.Records[0].cf.request;

  /**
   * This blueprint demonstrates how an origin-request trigger can be used to
   * change the origin from which the content is fetched, based on request
   properties.
   * In this example, we use the value of the CloudFront-Viewer-Country header
   * to update the S3 bucket domain name to a bucket in a Region that is closer to
   * the viewer.
   *
   * This can be useful in several ways:
   * 1) Reduces latencies when the Region specified is nearer to the viewer's
   *    country.
   * 2) Provides data sovereignty by making sure that data is served from an
   *    origin that's in the same country that the request came from.
   *
   * NOTE: 1. You must configure your distribution to cache based on the
   *         CloudFront-Viewer-Country header. For more information, see
   *         http://docs.aws.amazon.com/console/cloudfront/cache-on-selected-headers
   *        2. CloudFront adds the CloudFront-Viewer-Country header after the viewer
   *         request event. To use this example, you must create a trigger for the
   *         origin request event.
   */

  const countryToRegion = {
    'DE': 'eu-central-1',
    'IE': 'eu-west-1',
    'GB': 'eu-west-2',
    'FR': 'eu-west-3',
    'JP': 'ap-northeast-1',
    'IN': 'ap-south-1'
  };

  if (request.headers['cloudfront-viewer-country']) {
    const countryCode = request.headers['cloudfront-viewer-country'][0].value;
    const region = countryToRegion[countryCode];

    /**
     * If the viewer's country is not in the list you specify, the request
     * goes to the default S3 bucket you've configured.
     */
    if (region) {
      /**
       * If you've set up OAI, the bucket policy in the destination bucket
       * should allow the OAI GetObject operation, as configured by default
       * for an S3 origin with OAI. Another requirement with OAI is to provide

```



```
        * the Region so it can be used for the SIGV4 signature. Otherwise, the
        * Region is not required.
        */
        request.origin.s3.region = region;
        const domainName = `my-bucket-in-${region}.s3.amazonaws.com`;
        request.origin.s3.domainName = domainName;
        request.headers['host'] = [{ key: 'host', value: domainName }];
    }
}

callback(null, request);
};
```

Python

```
def lambda_handler(event, context):
    request = event['Records'][0]['cf']['request']

    '''
    This blueprint demonstrates how an origin-request trigger can be used to
    change the origin from which the content is fetched, based on request properties.
    In this example, we use the value of the CloudFront-Viewer-Country header
    to update the S3 bucket domain name to a bucket in a Region that is closer to
    the viewer.

    This can be useful in several ways:
    1) Reduces latencies when the Region specified is nearer to the viewer's
       country.
    2) Provides data sovereignty by making sure that data is served from an
       origin that's in the same country that the request came from.

    NOTE: 1. You must configure your distribution to cache based on the
           CloudFront-Viewer-Country header. For more information, see
           http://docs.aws.amazon.com/console/cloudfront/cache-on-selected-headers
           2. CloudFront adds the CloudFront-Viewer-Country header after the viewer
           request event. To use this example, you must create a trigger for the
           origin request event.
    '''

    countryToRegion = {
        'DE': 'eu-central-1',
        'IE': 'eu-west-1',
        'GB': 'eu-west-2',
        'FR': 'eu-west-3',
        'JP': 'ap-northeast-1',
        'IN': 'ap-south-1'
    }

    viewerCountry = request['headers'].get('cloudfront-viewer-country')
    if viewerCountry:
        countryCode = viewerCountry[0]['value']
        region = countryToRegion.get(countryCode)

        # If the viewer's country is not in the list you specify, the request
        # goes to the default S3 bucket you've configured
        if region:
            '''
            If you've set up OAI, the bucket policy in the destination bucket
            should allow the OAI GetObject operation, as configured by default
            for an S3 origin with OAI. Another requirement with OAI is to provide
            the Region so it can be used for the SIGV4 signature. Otherwise, the
            Region is not required.
            '''
            request['origin']['s3']['region'] = region
            domainName = 'my-bucket-in-%s.s3.amazonaws.com' % region
```

```
request['origin']['s3']['domainName'] = domainName
request['headers']['host'] = [{'key': 'host', 'value': domainName}]

return request
```

Example: Using an Origin-Request Trigger to Change From an Amazon S3 Origin to a Custom Origin

This function demonstrates how an origin-request trigger can be used to change the custom origin from which the content is fetched, based on request properties.

Node.js

```
'use strict';

const querystring = require('querystring');

exports.handler = (event, context, callback) => {
  const request = event.Records[0].cf.request;

  /**
   * Reads query string to check if custom origin should be used, and
   * if true, sets custom origin properties.
   */

  const params = querystring.parse(request.querystring);

  if (params['useCustomOrigin']) {
    if (params['useCustomOrigin'] === 'true') {

      /* Set custom origin fields*/
      request.origin = {
        custom: {
          domainName: 'www.example.com',
          port: 443,
          protocol: 'https',
          path: '',
          sslProtocols: ['TLSv1', 'TLSv1.1'],
          readTimeout: 5,
          keepaliveTimeout: 5,
          customHeaders: {}
        }
      };
      request.headers['host'] = [{ key: 'host', value: 'www.example.com'}];
    }
  }
  callback(null, request);
};
```

Python

```
from urllib.parse import parse_qs

def lambda_handler(event, context):
    request = event['Records'][0]['cf']['request']

    # Reads query string to check if custom origin should be used, and
    # if true, sets custom origin properties

    params = {k: v[0] for k, v in parse_qs(request['querystring']).items()}
```

```
if params.get('useCustomOrigin') == 'true':
    # Set custom origin fields
    request['origin'] = {
        'custom': {
            'domainName': 'www.example.com',
            'port': 443,
            'protocol': 'https',
            'path': '',
            'sslProtocols': ['TLSv1', 'TLSv1.1'],
            'readTimeout': 5,
            'keepaliveTimeout': 5,
            'customHeaders': {}
        }
    }
    request['headers']['host'] = [{'key': 'host', 'value': 'www.example.com'}]

return request
```

Example: Using an Origin-Request Trigger to Gradually Transfer Traffic From One Amazon S3 Bucket to Another

This function demonstrates how you can gradually transfer traffic from one Amazon S3 bucket to another, in a controlled way.

Node.js

```
'use strict';

function getRandomInt(min, max) {
    /* Random number is inclusive of min and max*/
    return Math.floor(Math.random() * (max - min + 1)) + min;
}

exports.handler = (event, context, callback) => {
    const request = event.Records[0].cf.request;
    const BLUE_TRAFFIC_PERCENTAGE = 80;

    /**
     * This Lambda function demonstrates how to gradually transfer traffic from
     * one S3 bucket to another in a controlled way.
     * We define a variable BLUE_TRAFFIC_PERCENTAGE which can take values from
     * 1 to 100. If the generated randomNumber less than or equal to
     * BLUE_TRAFFIC_PERCENTAGE, traffic
     * is re-directed to blue-bucket. If not, the default bucket that we've configured
     * is used.
     */

    const randomNumber = getRandomInt(1, 100);

    if (randomNumber <= BLUE_TRAFFIC_PERCENTAGE) {
        const domainName = 'blue-bucket.s3.amazonaws.com';
        request.origin.s3.domainName = domainName;
        request.headers['host'] = [{ key: 'host', value: domainName}];
    }
    callback(null, request);
};
```

Python

```
import math
import random
```

```
def getRandomInt(min, max):
    # Random number is inclusive of min and max
    return math.floor(random.random() * (max - min + 1)) + min

def lambda_handler(min, max):
    request = event['Records'][0]['cf']['request']
    BLUE_TRAFFIC_PERCENTAGE = 80

    '''
    This Lambda function demonstrates how to gradually transfer traffic from
    one S3 bucket to another in a controlled way.
    We define a variable BLUE_TRAFFIC_PERCENTAGE which can take values from
    1 to 100. If the generated randomNumber less than or equal to
    BLUE_TRAFFIC_PERCENTAGE, traffic
    is re-directed to blue-bucket. If not, the default bucket that we've configured
    is used.
    '''

    randomNumber = getRandomInt(1, 100)

    if randomNumber <= BLUE_TRAFFIC_PERCENTAGE:
        domainName = 'blue-bucket.s3.amazonaws.com'
        request['origin']['s3']['domainName'] = domainName
        request['headers']['host'] = [{key: 'host', 'value': domainName}]

    return request
```

Example: Using an Origin-Request Trigger to Change the Origin Domain Name Based on the Country Header

This function demonstrates how you can change the origin domain name based on the CloudFront-Viewer-Country header, so content is served from an origin closer to the viewer's country.

Implementing this functionality for your distribution can have advantages such as the following:

- Reducing latencies when the Region specified is nearer to the viewer's country
- Providing data sovereignty by making sure that data is served from an origin that's in the same country that the request came from

Note that to enable this functionality you must configure your distribution to cache based on the CloudFront-Viewer-Country header. For more information, see [Cache Based on Selected Request Headers \(p. 50\)](#).

Node.js

```
'use strict';

exports.handler = (event, context, callback) => {
    const request = event.Records[0].cf.request;

    if (request.headers['cloudfront-viewer-country']) {
        const countryCode = request.headers['cloudfront-viewer-country'][0].value;
        if (countryCode === 'GB' || countryCode === 'DE' || countryCode === 'IE' ) {
            const domainName = 'eu.example.com';
            request.origin.custom.domainName = domainName;
            request.headers['host'] = [{key: 'host', value: domainName}];
        }
    }
}
```

```
    callback(null, request);  
};
```

Python

```
def lambda_handler(event, context):  
    request = event['Records'][0]['cf']['request']  
  
    viewerCountry = request['headers'].get('cloudfront-viewer-country')  
    if viewerCountry:  
        countryCode = viewerCountry[0]['value']  
        if countryCode == 'GB' or countryCode == 'DE' or countryCode == 'IE':  
            domainName = 'eu.example.com'  
            request['origin']['custom']['domainName'] = domainName  
            request['headers']['host'] = [{'key': 'host', 'value': domainName}]  
    return request
```

Updating Error Statuses - Examples

The examples in this section provide guidance for how you can use Lambda@Edge to change the error status that is returned to users.

Topics

- [Example: Using an Origin-Response Trigger to Update the Error Status Code to 200-OK \(p. 388\)](#)
- [Example: Using an Origin-Response Trigger to Update the Error Status Code to 302-Found \(p. 389\)](#)

Example: Using an Origin-Response Trigger to Update the Error Status Code to 200-OK

This function demonstrates how you can update the response status to 200 and generate static body content to return to the viewer in the following scenario:

- The function is triggered in an origin response.
- The response status from the origin server is an error status code (4xx or 5xx).

Node.js

```
'use strict';  
  
exports.handler = (event, context, callback) => {  
    const response = event.Records[0].cf.response;  
  
    /**  
     * This function updates the response status to 200 and generates static  
     * body content to return to the viewer in the following scenario:  
     * 1. The function is triggered in an origin response  
     * 2. The response status from the origin server is an error status code (4xx or  
5xx)  
     */  
  
    if (response.status >= 400 && response.status <= 599) {  
        response.status = 200;  
        response.statusDescription = 'OK';  
        response.body = 'Body generation example';  
    }  
}
```

```
    callback(null, response);  
};
```

Python

```
def lambda_handler(event, context):  
    response = event['Records'][0]['cf']['response']  
  
    '''  
    This function updates the response status to 200 and generates static  
    body content to return to the viewer in the following scenario:  
    1. The function is triggered in an origin response  
    2. The response status from the origin server is an error status code (4xx or 5xx)  
    '''  
  
    if int(response['status']) >= 400 and int(response['status']) <= 599:  
        response['status'] = 200  
        response['statusDescription'] = 'OK'  
        response['body'] = 'Body generation example'  
    return response
```

Example: Using an Origin-Response Trigger to Update the Error Status Code to 302-Found

This function demonstrates how you can update the HTTP status code to 302 to redirect to another path (cache behavior) that has a different origin configured. Note the following:

- The function is triggered in an origin response.
- The response status from the origin server is an error status code (4xx or 5xx).

Node.js

```
'use strict';  
  
exports.handler = (event, context, callback) => {  
    const response = event.Records[0].cf.response;  
    const request = event.Records[0].cf.request;  
  
    /**  
     * This function updates the HTTP status code in the response to 302, to redirect  
     * to another  
     * path (cache behavior) that has a different origin configured. Note the  
     * following:  
     * 1. The function is triggered in an origin response  
     * 2. The response status from the origin server is an error status code (4xx or  
     * 5xx)  
     */  
  
    if (response.status >= 400 && response.status <= 599) {  
        const redirect_path = `/${request.cf.path}?${request.cf.querystring}`;  
  
        response.status = 302;  
        response.statusDescription = 'Found';  
  
        /* Drop the body, as it is not required for redirects */  
        response.body = '';  
        response.headers['location'] = [{ key: 'Location', value: redirect_path }];  
    }  
}
```

```
        callback(null, response);  
    };
```

Python

```
def lambda_handler(event, context):  
    response = event['Records'][0]['cf']['response']  
    request = event['Records'][0]['cf']['request']  
  
    '''  
    This function updates the HTTP status code in the response to 302, to redirect to  
    another  
    path (cache behavior) that has a different origin configured. Note the following:  
    1. The function is triggered in an origin response  
    2. The response status from the origin server is an error status code (4xx or 5xx)  
    '''  
  
    if int(response['status']) >= 400 and int(response['status']) <= 599:  
        redirect_path = '/plan-b/path?%s' % request['querystring']  
  
        response['status'] = 302  
        response['statusDescription'] = 'Found'  
  
        # Drop the body as it is not required for redirects  
        response['body'] = ''  
        response['headers']['location'] = [{'key': 'Location', 'value':  
redirect_path}]  
  
    return response
```

Accessing the Request Body - Examples

The examples in this section illustrate how you can use Lambda@Edge to work with POST requests.

Note

To use these examples, you must enable the *include body* option in the distribution's Lambda function association. It is not enabled by default.

- To enable this setting in the CloudFront console, select the check box for **Include Body** in the **Lambda Function Association**.
- To enable this setting in the CloudFront API or with AWS CloudFormation, set the `IncludeBody` field to `true` in `LambdaFunctionAssociation`.

Topics

- [Example: Using a Request Trigger to Read an HTML Form \(p. 390\)](#)
- [Example: Using a Request Trigger to Modify an HTML Form \(p. 392\)](#)

Example: Using a Request Trigger to Read an HTML Form

This function demonstrates how you can process the body of a POST request generated by an HTML form (web form), such as a "contact us" form. For example, you might have an HTML form like the following:

```
<html>  
  <form action="http://example.com" method="post">  
    Param 1: <input type="text" name="name1"><br>
```

```
Param 2: <input type="text" name="name2"><br>
input type="submit" value="Submit">
</form>
</html>
```

For the example function that follows, the function must be triggered in a CloudFront viewer request or origin request.

Node.js

```
'use strict';

const querystring = require('querystring');

/**
 * This function demonstrates how you can read the body of a POST request
 * generated by an HTML form (web form). The function is triggered in a
 * CloudFront viewer request or origin request event type.
 */

exports.handler = (event, context, callback) => {
  const request = event.Records[0].cf.request;

  if (request.method === 'POST') {
    /* HTTP body is always passed as base64-encoded string. Decode it. */
    const body = Buffer.from(request.body.data, 'base64').toString();

    /* HTML forms send the data in query string format. Parse it. */
    const params = querystring.parse(body);

    /* For demonstration purposes, we only log the form fields here.
     * You can put your custom logic here. For example, you can store the
     * fields in a database, such as AWS DynamoDB, and generate a response
     * right from your Lambda@Edge function.
     */
    for (let param in params) {
      console.log(`For "${param}" user submitted "${params[param]}".\n`);
    }
  }
  return callback(null, request);
};
```

Python

```
import base64
from urllib.parse import parse_qs

'''
Say there is a POST request body generated by an HTML such as:

<html>
  <form action="http://example.com" method="post">
    Param 1: <input type="text" name="name1"><br>
    Param 2: <input type="text" name="name2"><br>
    input type="submit" value="Submit">
  </form>
</html>

'''

'''
This function demonstrates how you can read the body of a POST request
generated by an HTML form (web form). The function is triggered in a
```



```
CloudFront viewer request or origin request event type.
'''

def lambda_handler(event, context):
    request = event['Records'][0]['cf']['request']

    if request['method'] == 'POST':
        # HTTP body is always passed as base64-encoded string. Decode it
        body = base64.b64decode(request['body']['data'])

        # HTML forms send the data in query string format. Parse it
        params = {k: v[0] for k, v in parse_qs(body).items()}

        '''
        For demonstration purposes, we only log the form fields here.
        You can put your custom logic here. For example, you can store the
        fields in a database, such as AWS DynamoDB, and generate a response
        right from your Lambda@Edge function.
        '''
        for key, value in params.items():
            print("For %s use submitted %s" % (key, value))

    return request
```

Example: Using a Request Trigger to Modify an HTML Form

This function demonstrates how you can modify the body of a POST request generated by an HTML form (web form). The function is triggered in a CloudFront viewer request or origin request.

Node.js

```
'use strict';

const querystring = require('querystring');

exports.handler = (event, context, callback) => {
    var request = event.Records[0].cf.request;
    if (request.method === 'POST') {
        /* Request body is being replaced. To do this, update the following
        /* three fields:
        *   1) body.action to 'replace'
        *   2) body.encoding to the encoding of the new data.
        *
        *       Set to one of the following values:
        *
        *       text - denotes that the generated body is in text format.
        *           Lambda@Edge will propagate this as is.
        *       base64 - denotes that the generated body is base64 encoded.
        *           Lambda@Edge will base64 decode the data before sending
        *           it to the origin.
        *   3) body.data to the new body.
        */
        request.body.action = 'replace';
        request.body.encoding = 'text';
        request.body.data = getUpdatedBody(request);
    }
    callback(null, request);
};

function getUpdatedBody(request) {
    /* HTTP body is always passed as base64-encoded string. Decode it. */
    const body = Buffer.from(request.body.data, 'base64').toString();
```

```
/* HTML forms send data in query string format. Parse it. */
const params = querystring.parse(body);

/* For demonstration purposes, we're adding one more param.
 *
 * You can put your custom logic here. For example, you can truncate long
 * bodies from malicious requests.
 */
params['new-param-name'] = 'new-param-value';
return querystring.stringify(params);
}
```

Python

```
import base64
from urllib.parse import parse_qs, urlencode

def lambda_handler(event, context):
    request = event['Records'][0]['cf']['request']
    if request['method'] == 'POST':
        '''
        Request body is being replaced. To do this, update the following
        three fields:
        1) body.action to 'replace'
        2) body.encoding to the encoding of the new data.

        Set to one of the following values:

            text - denotes that the generated body is in text format.
                Lambda@Edge will propagate this as is.
            base64 - denotes that the generated body is base64 encoded.
                Lambda@Edge will base64 decode the data before sending
                it to the origin.
        3) body.data to the new body.
        '''
        request['body']['action'] = 'replace'
        request['body']['encoding'] = 'text'
        request['body']['data'] = getUpdatedBody(request)
    return request

def getUpdatedBody(request):
    # HTTP body is always passed as base64-encoded string. Decode it
    body = base64.b64decode(request['body']['data'])

    # HTML forms send data in query string format. Parse it
    params = {k: v[0] for k, v in parse_qs(body).items()}

    # For demonstration purposes, we're adding one more param

    # You can put your custom logic here. For example, you can truncate long
    # bodies from malicious requests
    params['new-param-name'] = 'new-param-value'
    return urlencode(params)
```

Requirements and Restrictions on Lambda Functions

See the following sections for requirements and restrictions on using Lambda functions with CloudFront.

Topics

- [CloudFront Distributions and Associations \(p. 394\)](#)
- [CloudFront Triggers for Lambda Functions \(p. 394\)](#)
- [CloudWatch Logs \(p. 394\)](#)
- [Headers \(p. 394\)](#)
- [HTTP Status Codes \(p. 396\)](#)
- [Lambda Function Supported Runtimes and Configuration \(p. 397\)](#)
- [Quotas \(p. 397\)](#)
- [Microsoft Smooth Streaming \(p. 399\)](#)
- [Network Access \(p. 399\)](#)
- [Query String Parameters \(p. 399\)](#)
- [Maximum Size for Body with the Include Body Option \(p. 399\)](#)
- [Tagging \(p. 399\)](#)
- [URI \(p. 400\)](#)
- [URI and Query String Encoding \(p. 400\)](#)

CloudFront Distributions and Associations

- You cannot associate a Lambda function with a CloudFront distribution owned by another AWS account.

CloudFront Triggers for Lambda Functions

- You can add triggers only for a numbered version, not for `$LATEST` or for aliases.
- You can add triggers only for functions in the US East (N. Virginia) Region.
- To add triggers, the IAM execution role associated with your Lambda function must be assumable by the service principals `lambda.amazonaws.com` and `edgelambda.amazonaws.com`. For more information, see [Setting IAM Permissions and Roles for Lambda@Edge \(p. 329\)](#).

CloudWatch Logs

For information about Amazon CloudWatch Logs quotas (formerly known as limits), see [CloudWatch Logs quotas](#) in the *Amazon CloudWatch User Guide*.

Headers

Note the following requirements and restrictions on using headers with Lambda@Edge.

Topics

- [Blacklisted Headers \(p. 394\)](#)
- [Read-only Headers \(p. 395\)](#)
- [CloudFront-* Headers \(p. 396\)](#)

Blacklisted Headers

Blacklisted headers aren't exposed and can't be added by Lambda@Edge functions. If your Lambda function adds a blacklisted header, the request fails CloudFront validation. CloudFront returns HTTP status code 502 (Bad Gateway) to the viewer.

- Connection
- Expect
- Keep-Alive
- Proxy-Authenticate
- Proxy-Authorization
- Proxy-Connection
- Trailer
- Upgrade
- X-Accel-Buffering
- X-Accel-Charset
- X-Accel-Limit-Rate
- X-Accel-Redirect
- X-Amz-Cf-*
- X-Amzn-*
- X-Cache
- X-Edge-*
- X-Forwarded-Proto
- X-Real-IP

Read-only Headers

Read-only headers can be read but not edited. You can use them as input to CloudFront caching logic, and your Lambda function can read the header values, but it can't change the values. If your Lambda function adds or edits a read-only header, the request fails CloudFront validation. CloudFront returns HTTP status code 502 (Bad Gateway) to the viewer.

Read-only Headers for CloudFront Viewer Request Events

- Content-Length
- Host
- Transfer-Encoding
- Via

Read-only Headers for CloudFront Origin Request Events

- Accept-Encoding
- Content-Length
- If-Modified-Since
- If-None-Match
- If-Range
- If-Unmodified-Since
- Transfer-Encoding
- Via

Read-only Headers for CloudFront Origin Response Events

- Transfer-Encoding

- Via

Read-only Headers for CloudFront Viewer Response Events

- Content-Encoding
- Content-Length
- Transfer-Encoding
- Warning
- Via

CloudFront-* Headers

A Lambda function can read, edit, remove, or add any of the following headers.

- CloudFront-Forwarded-Proto
- CloudFront-Is-Desktop-Viewer
- CloudFront-Is-Mobile-Viewer
- CloudFront-Is-SmartTV-Viewer
- CloudFront-Is-Tablet-Viewer
- CloudFront-Viewer-Country

Note the following:

- If you want CloudFront to add these headers, you must configure CloudFront to cache based on these headers. For information about configuring CloudFront to cache based on specified headers, see [Cache Based on Selected Request Headers \(p. 50\)](#) in the topic [Values That You Specify When You Create or Update a Distribution \(p. 38\)](#).
- CloudFront adds the headers after the viewer request event.
- If the viewer adds headers that have these names, CloudFront overwrites the header values.
- For viewer events, `CloudFront-Viewer-Country` is blacklisted. Blacklisted headers aren't exposed and can't be added by Lambda@Edge functions. If your Lambda function adds a blacklisted header, the request fails CloudFront validation, and CloudFront returns HTTP status code 502 (Bad Gateway) to the viewer.

For more information, see the following examples:

- [Example: Redirecting Viewer Requests to a Country-Specific URL \(p. 378\)](#)
- [Example: Serving Different Versions of an Object Based on the Device \(p. 379\)](#)

HTTP Status Codes

CloudFront doesn't execute Lambda functions for viewer response events if the origin returns HTTP status code 400 or higher. You also can't modify the HTTP status code from a viewer response event for a successful request.

You can, however, execute Lambda functions for *origin* response errors, including HTTP status codes 4xx and 5xx. For more information, see [Updating HTTP Responses in Origin-Response Triggers \(p. 366\)](#).

Lambda Function Supported Runtimes and Configuration

- Lambda@Edge supports Lambda functions with the following runtimes:

- Python 3.8
- Python 3.7
- Node.js 12
- Node.js 10
- Node.js 8 and Node.js 6

Note

Node.js versions 8 and 6 have reached end of life. You can't create or update functions with these runtimes. If you have an existing function with one of these runtimes you can still associate it with a CloudFront distribution, and functions that are already associated with a distribution will still run. However, we recommend moving your function to a newer version of Node.js. For more information, see [Runtime Support Policy](#) in the *AWS Lambda Developer Guide* and the [Node.js release schedule on GitHub](#).

- You can't configure your Lambda function to access resources inside your VPC.
- You can't associate your Lambda function with a CloudFront distribution owned by another AWS account.
- [AWS Lambda function dead letter queues](#) are not supported.
- [AWS Lambda environment variables](#) are not supported.
- Functions with [AWS Lambda layers](#) are not supported.
- [Using AWS X-Ray](#) is not supported.
- [AWS Lambda reserved concurrency and provisioned concurrency](#) are not supported.

Quotas

The quotas in this section apply to Lambda@Edge. These quotas are in addition to the default CloudFront and Lambda quotas, which also apply. For the default quotas, see [Quotas \(p. 496\)](#) in this guide and [Quotas](#) in the *AWS Lambda Developer Guide*.

Note

Lambda dynamically scales capacity in response to increased traffic, within your account's quotas. For more information, see [Function Scaling](#) in the *AWS Lambda Developer Guide*.

In addition, be aware that there are some other restrictions when using Lambda@Edge functions. For more information, see [Requirements and Restrictions on Lambda Functions \(p. 393\)](#).

Quotas that differ by event type

Entity	Origin request and response event quotas	Viewer request and response event quotas
Function memory size	Same as Lambda quotas	128 MB
Function timeout. The function can make network calls to resources such as Amazon S3 buckets, DynamoDB tables, or Amazon EC2 instances in AWS Regions.	30 seconds	5 seconds

Entity	Origin request and response event quotas	Viewer request and response event quotas
Size of a response that is generated by a Lambda function, including headers and body	1 MB	40 KB
Maximum compressed size of a Lambda function and any included libraries	50 MB	1 MB

Other quotas

Entity	Default quota
Distributions per AWS account that you can create triggers for	25 Request a higher quota
Triggers per distribution	100 Request a higher quota
Requests per second	10,000 (in each Region) Request a higher quota
Concurrent executions	1000 (in each Region) Request a higher quota
For more information, see Function Scaling in the <i>AWS Lambda Developer Guide</i> .	

Size Quotas on URI and Query String

When accessing or updating a URI or query string in a Lambda@Edge function, the total length of the URI including the query string must be less than 8,192 characters.

Size Quotas on Request Body with the Include Body Option

When you choose the **Include Body** option to expose the request body to your Lambda@Edge function, the following size quotas apply to the portions of the body that are exposed or replaced.

Note

The body is always base64 encoded by Lambda@Edge before it is exposed.

Size Quotas when Exposing the Body to a Lambda Function

If the request body is large, Lambda@Edge truncates it before exposing it, as follows:

- For viewer requests, the body is truncated at 40 KB.
- For origin requests, the body is truncated at 1 MB.

Size Quotas when Returning a Request Body from a Lambda Function

If you access the request body as read-only, the full original request body is sent to the origin. However, if you choose to replace the request body, the following body size quotas apply when it's returned from a Lambda function:

Type of body encoding	Allowed body size: Viewer request	Allowed body size: Origin request
text	40 KB	1 MB
base64	53.2 KB	1.33 MB

Microsoft Smooth Streaming

You can't create triggers for a CloudFront distribution that you're using for video on demand (VOD) streaming of media files that you've transcoded into the Microsoft Smooth Streaming format.

Network Access

Functions triggered by origin request and response events as well as functions triggered by viewer request and response events can make network calls to resources on the internet, and to AWS services such as Amazon S3 buckets, DynamoDB tables, or Amazon EC2 instances.

Query String Parameters

- To access a query string in a Lambda function, use `event.Records[0].cf.request.querystring`.
- A function can update a query string for viewer and origin request events. The updated query string can't include spaces, control characters, or the fragment identifier (#).
- A function can read a query string only for origin and viewer response events.
- Configuring CloudFront to cache based on query string parameters affects whether a function can access the query string:
 - **Viewer request and response events** – A function can access a query string regardless of the setting for [Query String Forwarding and Caching \(p. 52\)](#).
 - **Origin request and response events** – A function can access a query string only if [Query String Forwarding and Caching \(p. 52\)](#) is set either to **Forward All, Cache Based on Whitelist** or to **Forward All, Cache Based on All**.
- We recommend that you use percent encoding for the URI and query string. For more information, see [URI and Query String Encoding \(p. 400\)](#).
- The total size of the URI (`event.Records[0].cf.request.uri`) and the query string (`event.Records[0].cf.request.querystring`) must be less than 8,192 characters.

For more information, see [Caching Content Based on Query String Parameters \(p. 241\)](#).

Maximum Size for Body with the Include Body Option

For all Lambda@Edge quotas (formerly known as limits), including size quotas, see [Quotas \(p. 397\)](#).

Tagging

Some AWS services, including Amazon CloudFront and AWS Lambda, support adding tags to resources within the service. However, at this time, you cannot apply tags to Lambda@Edge resources. To learn more about tagging in CloudFront, see [Tagging Amazon CloudFront Distributions \(p. 64\)](#).

URI

If a function changes the URI for a request, that doesn't change the cache behavior for the request or the origin that the request is forwarded to.

URI and Query String Encoding

Lambda functions require the URI and query string to be UTF-8 encoded. (Percent encoding is compatible with UTF-8 encoding.) The behavior of CloudFront and Lambda depends on the following:

- The encoding of the URI and query string that CloudFront received in the request from the viewer
- Whether the URI or query string is changed by a function that is triggered by a viewer request or origin request event

Values Are UTF-8 Encoded

CloudFront forwards the values to your Lambda function without changing them.

Values Are ISO 8859-1 Encoded

CloudFront converts [ISO 8859-1 character encoding](#) to UTF-8 encoding before forwarding the values to your Lambda function.

Values Are Encoded Using Some Other Character Encoding

If the values are encoded using any other character encoding, CloudFront assumes that they're ISO 8859-1 encoded and tries to convert from ISO 8859-1 encoding to UTF-8 encoding.

Important

The converted version might be an inaccurate interpretation of the values in the original request. This can cause a Lambda function or your origin to produce an unintended result.

The value that CloudFront forwards to your origin server depends on whether functions that are triggered by viewer request or origin request events change the URI or query string:

- **If the functions don't change the URI or query string** – CloudFront forwards the values that CloudFront received in the request from the viewer to your origin server.
- **If the functions do change the URI or query string** – CloudFront forwards the UTF-8 encoded value.

In both cases, the behavior is unaffected by the character encoding of the request from the viewer.

Reports, Metrics, and Logs

This section includes topics that provide details on your options for reports and monitoring for CloudFront. A variety of reports are available for you to see usage and activity for your CloudFront distributions, including billing reports, cache statistics, popular content, and top referrers. In addition, you can monitor and track CloudFront—including Lambda@Edge activity—directly in the CloudFront console, and by using tools such as CloudTrail and CloudWatch.

Topics

- [AWS Billing and Usage Reports for CloudFront \(p. 401\)](#)
- [CloudFront Reports in the Console \(p. 405\)](#)
- [Monitoring CloudFront with Amazon CloudWatch \(p. 428\)](#)
- [CloudFront logging \(p. 438\)](#)
- [Tracking Configuration Changes with AWS Config \(p. 472\)](#)

AWS Billing and Usage Reports for CloudFront

AWS provides two usage reports for CloudFront:

- The billing report is a high-level view of all of the activity for the AWS services that you're using, including CloudFront. For more information, see [AWS Billing Report for CloudFront \(p. 401\)](#).
- The usage report is a summary of activity for a specific service, aggregated by hour, day, or month. It also includes usage charts that provide a graphical representation of your CloudFront usage. For more information, see [AWS Usage Report for CloudFront \(p. 402\)](#).

To help you understand these reports, see the detailed information in [Interpreting Your AWS Bill and the AWS Usage Report for CloudFront \(p. 403\)](#).

Note

Like other AWS services, CloudFront charges you for only what you use. For more information, see [CloudFront pricing \(p. 8\)](#).

Topics

- [AWS Billing Report for CloudFront \(p. 401\)](#)
- [AWS Usage Report for CloudFront \(p. 402\)](#)
- [Interpreting Your AWS Bill and the AWS Usage Report for CloudFront \(p. 403\)](#)

AWS Billing Report for CloudFront

You can view a summary of your AWS usage and charges, listed by service, on the Bills page in the AWS Management Console.

You can also download a more detailed version of the report in CSV format. The detailed billing report includes the following values that apply to CloudFront:

- **ProductCode** — AmazonCloudFront

- **UsageType** — One of the following values
 - A code that identifies the type of data transfer
 - Invalidations
 - SSL-Cert-Custom

For more information, see [Interpreting Your AWS Bill and the AWS Usage Report for CloudFront \(p. 403\)](#).

- **ItemDescription** — A description of the billing rate for the **UsageType**.
- **Usage Start Date/Usage End Date** — The day that the usage applies to, in Coordinated Universal Time (UTC).
- **Usage Quantity** — One of the following values:
 - The number of requests during the specified time period
 - The amount of data transferred in gigabytes
 - The number of objects invalidated
 - The sum of the prorated months that you had SSL certificates associated with enabled CloudFront distributions. For example, if you have one certificate associated with an enabled distribution for an entire month and another certificate associated with an enabled distribution for half of the month, this value will be 1.5.

To display summary billing information and download the detailed billing report

1. Sign in to the AWS Management Console at <https://console.aws.amazon.com/console/home>.
2. In the title bar, click your IAM user name, and click **Billing & Cost Management**.
3. In the navigation pane, click **Bills**.
4. To view summary information for CloudFront, under **Details**, click **CloudFront**.
5. To download a detailed billing report in CSV format, click **Download CSV**, and follow the on-screen prompts to save the report.

AWS Usage Report for CloudFront

AWS provides a CloudFront usage report that is more detailed than the billing report but less detailed than CloudFront access logs. The usage report provides aggregate usage data by hour, day, or month; and it lists operations by region and usage type, such as data transferred out of the Australia region.

The CloudFront usage report includes the following values:

- **Service** — AmazonCloudFront
- **Operation** — HTTP method. Values include DELETE, GET, HEAD, OPTIONS, PATCH, POST, and PUT.
- **UsageType** — One of the following values
 - A code that identifies the type of data transfer
 - Invalidations
 - SSL-Cert-Custom

For more information, see [Interpreting Your AWS Bill and the AWS Usage Report for CloudFront \(p. 403\)](#).

- **Resource** — Either the ID of the CloudFront distribution associated with the usage or the certificate ID of an SSL certificate that you have associated with a CloudFront distribution.
- **StartTime/EndTime** — The day that the usage applies to, in Coordinated Universal Time (UTC).
- **UsageValue** — (1) The number of requests during the specified time period or (2) the amount of data transferred in bytes.

If you're using Amazon S3 as the origin for CloudFront, consider running the usage report for Amazon S3, too. However, if you use Amazon S3 for purposes other than as an origin for your CloudFront distributions, it might not be clear what portion applies to your CloudFront usage.

Tip

For detailed information about every request that CloudFront receives for your objects, turn on CloudFront access logs for your distribution. For more information, see [Configuring and Using Standard Logs \(Access Logs\)](#) (p. 439).

To download the usage report for CloudFront or Amazon S3

1. Sign in to the AWS Management Console at <https://console.aws.amazon.com/console/home>.
2. In the title bar, click your IAM user name, and click **Billing & Cost Management**.
3. In the navigation pane, click **Reports**.
4. Under **AWS Usage Report**, click **AWS Usage Report**.
5. In the **Service** list, click **CloudFront** or **Amazon Simple Storage Service**.
6. Select the settings to use:
 - **Usage Types** — For a detailed explanation of CloudFront usage types, see [the section called “Interpreting Your AWS Bill and the AWS Usage Report for CloudFront”](#) (p. 403).

For Amazon S3, select **All Usage Types**.

 - **Operation** — Select **All Operations**.
 - **Time Period** — Select the time period that you want the report to cover.
 - **Report Granularity** — Select whether you want the report to include subtotals by the hour, by the day, or by the month.
7. Click the download button for the desired format.
8. Follow the on-screen prompts to view or save the report.

Interpreting Your AWS Bill and the AWS Usage Report for CloudFront

Your AWS bill for CloudFront service includes codes and abbreviations that might not be immediately obvious. The first column in the following table lists items that appear in your bill and explains what each means.

In addition, you can get an AWS usage report for CloudFront that contains more detail than the AWS bill for CloudFront. The second column in the table lists items that appear in the usage report and shows the correlation between bill items and usage report items.

Most codes in both columns include a two-letter abbreviation that indicates the location of the activity. In the following table, *region* in a code is replaced in your AWS bill and in the usage report by one of the following two-letter abbreviations :

- **AP:** Hong Kong, Philippines, South Korea, Taiwan, and Singapore (Asia Pacific)
- **AU:** Australia
- **CA:** Canada
- **EU:** Europe and Israel
- **IN:** India
- **JP:** Japan
- **ME:** Middle East

- **SA:** South America
- **US:** United States
- **ZA:** South Africa

For more information about pricing by region, see [Amazon CloudFront Pricing](#).

Note

This table doesn't include charges for transferring your objects from an Amazon S3 bucket to CloudFront edge locations. These charges, if any, appear in the **AWS Data Transfer** portion of your AWS bill.

Items in Your CloudFront Bill	Values in the Usage Type Column in the CloudFront Usage Report
<p><i>region</i>-DataTransfer-Out-Bytes</p> <p>Sum of bytes that CloudFront served for web and RTMP distributions:</p> <ul style="list-style-type: none"> • Web distributions: Total bytes served from CloudFront edge locations in <i>region</i> in response to user <code>GET</code> and <code>HEAD</code> requests • RTMP distributions: Total bytes transferred from CloudFront edge locations in <i>region</i> to end users 	<p>Web distributions:</p> <ul style="list-style-type: none"> • <i>region</i>-Out-Bytes-HTTP-Static: Bytes served via HTTP for objects with TTL \geq 3600 seconds • <i>region</i>-Out-Bytes-HTTPS-Static: Bytes served via HTTPS for objects with TTL \geq 3600 seconds • <i>region</i>-Out-Bytes-HTTP-Dynamic: Bytes served via HTTP for objects with TTL $<$ 3600 seconds • <i>region</i>-Out-Bytes-HTTPS-Dynamic: Bytes served via HTTPS for objects with TTL $<$ 3600 seconds • <i>region</i>-Out-Bytes-HTTP-Proxy: Bytes returned from CloudFront to viewers via HTTP in response to <code>DELETE</code>, <code>OPTIONS</code>, <code>PATCH</code>, <code>POST</code>, and <code>PUT</code> requests. • <i>region</i>-Out-Bytes-HTTPS-Proxy: Bytes returned from CloudFront to viewers via HTTPS in response to <code>DELETE</code>, <code>OPTIONS</code>, <code>PATCH</code>, <code>POST</code>, and <code>PUT</code> requests. <p>RTMP distributions:</p> <ul style="list-style-type: none"> • <i>region</i>-FMS-Out-Bytes
<p><i>region</i>-DataTransfer-Out-OBytes</p> <p>Web distributions only: Total bytes transferred from CloudFront edge locations to your origin in response to <code>DELETE</code>, <code>OPTIONS</code>, <code>PATCH</code>, <code>POST</code>, and <code>PUT</code> requests. The charges include data transfer for WebSocket data from client to server.</p>	<p><i>region</i>-Out-OBytes-HTTP-Proxy</p> <p>Total bytes transferred via HTTP from CloudFront edge locations to your origin in response to <code>DELETE</code>, <code>OPTIONS</code>, <code>PATCH</code>, <code>POST</code>, and <code>PUT</code> requests.</p> <p><i>region</i>-Out-OBytes-HTTPS-Proxy</p> <p>Total bytes transferred via HTTPS from CloudFront edge locations to your origin in response to <code>DELETE</code>, <code>OPTIONS</code>, <code>PATCH</code>, <code>POST</code>, and <code>PUT</code> requests.</p>
<p><i>region</i>-Requests-Tier1</p> <p>Web distributions only: Number of HTTP <code>GET</code> and <code>HEAD</code> requests</p>	<p><i>region</i>-Requests-HTTP-Static</p> <p>Number of HTTP <code>GET</code> and <code>HEAD</code> requests served for objects with TTL \geq 3600 seconds</p> <p><i>region</i>-Requests-HTTP-Dynamic</p> <p>Number of HTTP <code>GET</code> and <code>HEAD</code> requests served for objects with TTL $<$ 3600 seconds</p>

Items in Your CloudFront Bill	Values in the Usage Type Column in the CloudFront Usage Report
<p><i>region</i>-Requests-Tier2-HTTPS</p> <p>Web distributions only: Number of HTTPS GET and HEAD requests</p>	<p><i>region</i>-Requests-HTTPS-Static</p> <p>Number of HTTPS GET and HEAD requests served for objects with TTL ≥ 3600 seconds</p> <p><i>region</i>-Requests-HTTPS-Dynamic</p> <p>Number of HTTPS GET and HEAD requests served for objects with TTL < 3600 seconds</p>
<p><i>region</i>-Requests-HTTP-Proxy</p> <p>Web distributions only: Number of HTTP DELETE, OPTIONS, PATCH, POST, and PUT requests that CloudFront forwards to your origin</p>	<p><i>region</i>-Requests-HTTP-Proxy</p> <p>Same as the corresponding item in your CloudFront bill</p>
<p><i>region</i>-Requests-HTTPS-Proxy</p> <p>Web distributions only: Number of HTTPS DELETE, OPTIONS, PATCH, POST, and PUT requests that CloudFront forwards to your origin</p>	<p><i>region</i>-Requests-HTTPS-Proxy</p> <p>Same as the corresponding item in your CloudFront bill</p>
<p><i>region</i>-Requests-HTTPS-Proxy-FLE</p> <p>Web distributions only: Number of HTTPS DELETE, OPTIONS, PATCH, and POST requests that CloudFront forwards to your origin which were processed with field-level encryption.</p>	<p><i>region</i>-Requests-HTTPS-Proxy-FLE</p> <p>Same as the corresponding item in your CloudFront bill</p>
<p>Invalidations</p> <p>Web distributions only: The charge for invalidating objects (removing the objects from CloudFront edge locations); for more information, see Paying for File Invalidation (p. 116).</p>	<p>Invalidations</p> <p>Same as the corresponding item in your CloudFront bill</p>
<p>SSL-Cert-Custom</p> <p>Web distributions only: The charge for using an SSL certificate with a CloudFront alternate domain name such as example.com instead of using the default CloudFront SSL certificate and the domain name that CloudFront assigned to your distribution.</p>	<p>SSL-Cert-Custom</p> <p>Same as the corresponding item in your CloudFront bill</p>

CloudFront Reports in the Console

The CloudFront console includes a variety of reports about your CloudFront activity, including the following:

- [CloudFront Cache Statistics Reports](#) (p. 406)
- [CloudFront Popular Objects Report](#) (p. 406)
- [CloudFront Top Referrers Report](#) (p. 406)
- [CloudFront Usage Reports](#) (p. 406)
- [CloudFront Viewers Reports](#) (p. 407)

Most of these reports are based on the data in CloudFront access logs, which contain detailed information about every user request that CloudFront receives. You don't need to enable access logs to view the reports. For more information, see [Configuring and Using Standard Logs \(Access Logs\)](#) (p. 439). The CloudFront usage report is based on the AWS usage report for CloudFront, which also doesn't require any special configuration. For more information, see [AWS Usage Report for CloudFront](#) (p. 402).

CloudFront Cache Statistics Reports

The CloudFront cache statistics report includes the following information:

- **Total Requests** – Shows the total number of requests for all HTTP status codes (for example, 200 or 404) and all methods (for example, GET, HEAD, or POST).
- **Percentage of Viewer Requests by Result Type** – Shows hits, misses, and errors as a percentage of total viewer requests for the selected CloudFront distribution.
- **Bytes Transferred to Viewers** – Shows total bytes and bytes from misses.
- **HTTP Status Codes** – Shows viewer requests by HTTP status code.
- **Percentage of GET Requests that Didn't Finish Downloading** – Shows viewer GET requests that didn't finish downloading the requested object as a percentage of total requests.

For more information, see [CloudFront Cache Statistics Reports](#) (p. 407).

CloudFront Popular Objects Report

The CloudFront popular objects report lists the 50 most popular objects and statistics about those objects, including the number of requests for the object, the number of hits and misses, the hit ratio, the number of bytes served for misses, the total bytes served, the number of incomplete downloads, and the number of requests by HTTP status code (2xx, 3xx, 4xx, and 5xx).

For more information, see [CloudFront Popular Objects Report](#) (p. 411).

CloudFront Top Referrers Report

The CloudFront top referrers report includes the top 25 referrers, the number of requests from a referrer, and the number of requests from a referrer as a percentage of the total number of requests during the specified period.

For more information, see [CloudFront Top Referrers Report](#) (p. 414).

CloudFront Usage Reports

The CloudFront usage reports include the following information:

- **Number of Requests** – Shows the total number of requests that CloudFront responds to from edge locations in the selected region during each time interval for the specified CloudFront distribution.
- **Data Transferred by Protocol and Data Transferred by Destination** – Both show the total amount of data transferred from CloudFront edge locations in the selected region during each time interval for the specified CloudFront distribution. They separate the data differently, as follows:
 - **By Protocol** – Separates the data by protocol: HTTP or HTTPS.
 - **By Destination** – Separates the data by destination: to your users or to your origin.

For more information, see [CloudFront Usage Reports \(p. 416\)](#).

CloudFront Viewers Reports

The CloudFront viewers reports include the following information:

- **Devices** – Shows the types of devices (for example, Desktop or Mobile) that your users use to access your content
- **Browsers** – Shows the name (or the name and version) of the browsers that your users use most frequently to access your content, for example, Chrome or Firefox
- **Operating Systems** – Shows the name (or the name and version) of the operating system that viewers run on most frequently when accessing your content, for example, Linux, Mac OS X, or Windows
- **Locations** – Shows the locations, by country or by U.S. state/territory, of the viewers that access your content most frequently

For more information, see [CloudFront Viewers Reports \(p. 421\)](#).

CloudFront Cache Statistics Reports

You can use the Amazon CloudFront console to display a graphical representation of statistics related to CloudFront edge locations. Data for these statistics are drawn from the same source as CloudFront access logs. You can display charts for a specified date range in the last 60 days, with data points every hour or every day. You can usually view data about requests that CloudFront received as recently as an hour ago, but data can occasionally be delayed by as much as 24 hours.

Note

You don't need to enable access logging to view cache statistics.

To display CloudFront cache statistics

1. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
2. In the navigation pane, click **Cache Statistics**.
3. In the **CloudFront Cache Statistics Reports** pane, for **Start Date** and **End Date**, select the date range for which you want to display cache statistics charts. Available ranges depend on the value that you select for **Granularity**:
 - **Daily** – To display charts with one data point per day, select any date range in the previous 60 days.
 - **Hourly** – To display charts with one data point every hour, select any date range of up to 14 days within the previous 60 days.

Dates and times are in Coordinated Universal Time (UTC).

4. For **Granularity**, specify whether to display one data point per day or one data point per hour in the charts. If you specify a date range greater than 14 days, the option to specify one data point per hour is not available.
5. For **Viewer Location**, choose the continent from which viewer requests originated, or choose **All Locations**. Cache statistics charts include data for requests that CloudFront received from the specified location.
6. In the **Distribution** list, select the distributions for which you want to display data in the usage charts:
 - **An individual web distribution** – The charts display data for the selected CloudFront web distribution. The **Distribution** list displays the distribution ID and alternate domain names

(CNAMEs) for the distribution, if any. If a distribution has no alternate domain names, the list includes origin domain names for the distribution.

- **All Web Distributions** – The charts display summed data for all web distributions that are associated with the current AWS account, excluding web distributions that you have deleted.
7. Click **Update**.
 8. To view data for a daily or hourly data point within a chart, move your mouse pointer over the data point.
 9. For charts that show data transferred, note that you can change the vertical scale to gigabytes, megabytes, or kilobytes for each chart.

Topics

- [Downloading Data in CSV Format \(p. 408\)](#)
- [How Cache Statistics Charts Are Related to Data in the CloudFront Access Logs \(p. 410\)](#)

Downloading Data in CSV Format

You can download the Cache Statistics report in CSV format. This section explains how to download the report and describes the values in the report.

To download the Cache Statistics report in CSV format

1. While viewing the Cache Statistics report, click **CSV**.
2. In the **Opening file name** dialog box, choose whether to open or save the file.

Information About the Report

The first few rows of the report include the following information:

Version

The version of the format for this CSV file.

Report

The name of the report.

DistributionID

The ID of the distribution that you ran the report for, or `ALL` if you ran the report for all distributions.

StartDateUTC

The beginning of the date range for which you ran the report, in Coordinated Universal Time (UTC).

EndDateUTC

The end of the date range for which you ran the report, in Coordinated Universal Time (UTC).

GeneratedTimeUTC

The date and time on which you ran the report, in Coordinated Universal Time (UTC).

Granularity

Whether each row in the report represents one hour or one day.

ViewerLocation

The continent that viewer requests originated from, or `ALL`, if you chose to download the report for all locations.

Data in the Cache Statistics Report

The report includes the following values:

DistributionID

The ID of the distribution that you ran the report for, or `ALL` if you ran the report for all distributions.

FriendlyName

An alternate domain name (CNAME) for the distribution, if any. If a distribution has no alternate domain names, the list includes an origin domain name for the distribution.

ViewerLocation

The continent that viewer requests originated from, or `ALL`, if you chose to download the report for all locations.

TimeBucket

The hour or the day that data applies to, in Coordinated Universal Time (UTC).

RequestCount

The total number of requests for all HTTP status codes (for example, 200 or 404) and all methods (for example, GET, HEAD, or POST).

HitCount

The number of viewer requests for which the object is served from a CloudFront edge cache.

MissCount

The number of viewer requests for which the object isn't currently in an edge cache, so CloudFront must get the object from your origin.

ErrorCount

The number of viewer requests that resulted in an error, so CloudFront didn't serve the object.

IncompleteDownloadCount

The number of viewer requests for which the viewer started but didn't finish downloading the object.

HTTP2xx

The number of viewer requests for which the HTTP status code was a 2xx value (succeeded).

HTTP3xx

The number of viewer requests for which the HTTP status code was a 3xx value (additional action is required).

HTTP4xx

The number of viewer requests for which the HTTP status code was a 4xx value (client error).

HTTP5xx

The number of viewer requests for which the HTTP status code was a 5xx value (server error).

TotalBytes

The total number of bytes served to viewers by CloudFront in response to all requests for all HTTP methods.

BytesFromMisses

The number of bytes served to viewers for objects that were not in the edge cache at the time of the request. This value is a good approximation of bytes transferred from your origin to CloudFront edge

caches. However, it excludes requests for objects that are already in the edge cache but that have expired.

How Cache Statistics Charts Are Related to Data in the CloudFront Access Logs

The following table shows how cache statistics charts in the CloudFront console correspond with values in CloudFront access logs. For more information about CloudFront access logs, see [Configuring and Using Standard Logs \(Access Logs\)](#) (p. 439).

Total Requests

This chart shows the total number of requests for all HTTP status codes (for example, 200 or 404) and all methods (for example, GET, HEAD, or POST). Total requests shown in this chart equal the total number of requests in the access log files for the same time period.

Percentage of Viewer Requests by Result Type

This chart shows hits, misses, and errors as a percentage of total viewer requests for the selected CloudFront distribution:

- **Hit** – A viewer request for which the object is served from a CloudFront edge cache. In access logs, these are requests for which the value of `x-edge-response-result-type` is `Hit`.
- **Miss** – A viewer request for which the object isn't currently in an edge cache, so CloudFront must get the object from your origin. In access logs, these are requests for which the value of `x-edge-response-result-type` is `Miss`.
- **Error** – A viewer request that resulted in an error, so CloudFront didn't serve the object. In access logs, these are requests for which the value of `x-edge-response-result-type` is `Error`, `LimitExceeded`, or `CapacityExceeded`.

The chart does not include refresh hits—requests for objects that are in the edge cache but that have expired. In access logs, refresh hits are requests for which the value of `x-edge-response-result-type` is `RefreshHit`.

Bytes Transferred to Viewers

This chart shows two values:

- **Total Bytes** – The total number of bytes served to viewers by CloudFront in response to all requests for all HTTP methods. In CloudFront access logs, **Total Bytes** is the sum of the values in the `sc-bytes` column for all of the requests during the same time period.
- **Bytes from Misses** – The number of bytes served to viewers for objects that were not in the edge cache at the time of the request. In CloudFront access logs, **Bytes from Misses** is the sum of the values in the `sc-bytes` column for requests for which the value of `x-edge-result-type` is `Miss`. This value is a good approximation of bytes transferred from your origin to CloudFront edge caches. However, it excludes requests for objects that are already in the edge cache but that have expired.

HTTP Status Codes

This chart shows viewer requests by HTTP status code. In CloudFront access logs, status codes appear in the `sc-status` column:

- **2xx** – The request succeeded.
- **3xx** – Additional action is required. For example, 301 (Moved Permanently) means that the requested object has moved to a different location.
- **4xx** – The client apparently made an error. For example, 404 (Not Found) means that the client requested an object that could not be found.
- **5xx** – The origin server didn't fill the request. For example, 503 (Service Unavailable) means that the origin server is currently unavailable.

Percentage of GET Requests that Didn't Finish Downloading

This chart shows viewer `GET` requests that didn't finish downloading the requested object as a percentage of total requests. Typically, downloading an object doesn't complete because the viewer canceled the download, for example, by clicking a different link or by closing the browser. In CloudFront access logs, these requests have a value of 200 in the `sc-status` column and a value of `Error` in the `x-edge-result-type` column.

CloudFront Popular Objects Report

The Amazon CloudFront console can display a list of the 50 most popular objects for a distribution during a specified date range in the previous 60 days.

Data for the Popular Objects report is drawn from the same source as CloudFront access logs. To get an accurate count of the top 50 objects, CloudFront counts the requests for all of your objects in 10-minute intervals beginning at midnight and keeps a running total of the top 150 objects for the next 24 hours. (CloudFront also retains daily totals for the top 150 objects for 60 days.) Near the bottom of the list, objects constantly rise onto or drop off of the list, so the totals for those objects are approximations. The fifty objects at the top of the list of 150 objects may rise and fall within the list, but they rarely drop off of the list altogether, so the totals for those objects typically are more reliable.

When an object drops off of the list of the top 150 objects and then rises onto the list again over the course of a day, CloudFront adds an estimated number of requests for the period that the object was missing from the list. The estimate is based on the number of requests received by whichever object was at the bottom of the list during that time period. If the object rises into the top 50 objects later in the day, the estimates of the number of requests that CloudFront received while the object was out of the top 150 objects usually causes the number of requests in the Popular Objects report to exceed the number of requests that appear in the access logs for that object.

Note

You don't need to enable access logging to view a list of popular objects.

To display popular objects for a distribution

1. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
2. In the navigation pane, click **Popular Objects**.
3. In the **CloudFront Popular Objects Report** pane, for **Start Date** and **End Date**, select the date range for which you want to display a list of popular objects. You can choose any date range in the previous 60 days.

Dates and times are in Coordinated Universal Time (UTC).

4. In the **Distribution** list, select the distribution for which you want to display a list of popular objects.
5. Click **Update**.

Topics

- [Downloading Data in CSV Format \(p. 411\)](#)
- [How Data in the Popular Objects Report Is Related to Data in the CloudFront Access Logs \(p. 413\)](#)

Downloading Data in CSV Format

You can download the Popular Objects report in CSV format. This section explains how to download the report and describes the values in the report.

To download the Popular Objects report in CSV format

1. While viewing the Popular Objects report, click **CSV**.
2. In the **Opening *file name*** dialog box, choose whether to open or save the file.

Information About the Report

The first few rows of the report include the following information:

Version

The version of the format for this CSV file.

Report

The name of the report.

DistributionID

The ID of the distribution that you ran the report for.

StartDateUTC

The beginning of the date range for which you ran the report, in Coordinated Universal Time (UTC).

EndDateUTC

The end of the date range for which you ran the report, in Coordinated Universal Time (UTC).

GeneratedTimeUTC

The date and time on which you ran the report, in Coordinated Universal Time (UTC).

Data in the Popular Objects Report

The report includes the following values:

DistributionID

The ID of the distribution that you ran the report for.

FriendlyName

An alternate domain name (CNAME) for the distribution, if any. If a distribution has no alternate domain names, the list includes an origin domain name for the distribution.

Object

The last 500 characters of the URL for the object.

RequestCount

The total number of requests for this object.

HitCount

The number of viewer requests for which the object is served from a CloudFront edge cache.

MissCount

The number of viewer requests for which the object isn't currently in an edge cache, so CloudFront must get the object from your origin.

HitCountPct

The value of `HitCount` as a percentage of the value of `RequestCount`.

BytesFromMisses

The number of bytes served to viewers for this object when the object was not in the edge cache at the time of the request.

TotalBytes

The total number of bytes served to viewers by CloudFront for this object in response to all requests for all HTTP methods.

IncompleteDownloadCount

The number of viewer requests for this object for which the viewer started but didn't finish downloading the object.

HTTP2xx

The number of viewer requests for which the HTTP status code was a 2xx value (succeeded).

HTTP3xx

The number of viewer requests for which the HTTP status code was a 3xx value (additional action is required).

HTTP4xx

The number of viewer requests for which the HTTP status code was a 4xx value (client error).

HTTP5xx

The number of viewer requests for which the HTTP status code was a 5xx value (server error).

How Data in the Popular Objects Report Is Related to Data in the CloudFront Access Logs

The following list shows how values in the Popular Objects report in the CloudFront console correspond with values in CloudFront access logs. For more information about CloudFront access logs, see [Configuring and Using Standard Logs \(Access Logs\)](#) (p. 439).

URL

The last 500 characters of the URL that viewers use to access the object.

Requests

The total number of requests for the object. This value generally corresponds closely with the number of `GET` requests for the object in CloudFront access logs.

Hits

The number of viewer requests for which the object was served from a CloudFront edge cache. In access logs, these are requests for which the value of `x-edge-response-result-type` is `Hit`.

Misses

The number of viewer requests for which the object wasn't in an edge cache, so CloudFront retrieved the object from your origin. In access logs, these are requests for which the value of `x-edge-response-result-type` is `Miss`.

Hit Ratio

The value of the **Hits** column as a percentage of the value of the **Requests** column.

Bytes from Misses

The number of bytes served to viewers for objects that were not in the edge cache at the time of the request. In CloudFront access logs, **Bytes from Misses** is the sum of the values in the `sc-bytes` column for requests for which the value of `x-edge-result-type` is `Miss`.

Total Bytes

The total number of bytes that CloudFront served to viewers in response to all requests for the object for all HTTP methods. In CloudFront access logs, **Total Bytes** is the sum of the values in the `sc-bytes` column for all of the requests during the same time period.

Incomplete Downloads

The number of viewer requests that did not finish downloading the requested object. Typically, the reason that a download doesn't complete is that the viewer canceled it, for example, by clicking a different link or by closing the browser. In CloudFront access logs, these requests have a value of 200 in the `sc-status` column and a value of `Error` in the `x-edge-result-type` column.

2xx

The number of requests for which the HTTP status code is 2xx, `Successful`. In CloudFront access logs, status codes appear in the `sc-status` column.

3xx

The number of requests for which the HTTP status code is 3xx, `Redirection`. 3xx status codes indicate that additional action is required. For example, 301 (Moved Permanently) means that the requested object has moved to a different location.

4xx

The number of requests for which the HTTP status code is 4xx, `Client Error`. 4xx status codes indicate that the client apparently made an error. For example, 404 (Not Found) means that the client requested an object that could not be found.

5xx

The number of requests for which the HTTP status code is 5xx, `Server Error`. 5xx status codes indicate that the origin server didn't fill the request. For example, 503 (Service Unavailable) means that the origin server is currently unavailable.

CloudFront Top Referrers Report

The CloudFront console can display a list of the 25 domains of the websites that originated the most HTTP and HTTPS requests for objects that CloudFront is distributing for a specified distribution. These top referrers can be search engines, other websites that link directly to your objects, or your own website. For example, if `http://example.com/index.html` links to 10 graphics, `example.com` is the referrer for all 10 graphics. You can display the Top Referrers report for any date range in the previous 60 days.

Note

If a user enters a URL directly into the address line of a browser, there is no referrer for the requested object.

Data for the Top Referrers report is drawn from the same source as CloudFront access logs. To get an accurate count of the top 25 referrers, CloudFront counts the requests for all of your objects in 10-minute intervals and keeps a running total of the top 75 referrers. Near the bottom of the list, referrers constantly rise onto or drop off of the list, so the totals for those referrers are approximations. The 25 referrers at the top of the list of 75 referrers may rise and fall within the list, but they rarely drop off of the list altogether, so the totals for those referrers typically are more reliable.

Note

You don't need to enable access logging to view a list of top referrers.

To display top referrers for a distribution

1. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
2. In the navigation pane, click **Top Referrers**.
3. In the **CloudFront Top Referrers Report** pane, for **Start Date** and **End Date**, select the date range for which you want to display a list of top referrers.

Dates and times are in Coordinated Universal Time (UTC).

4. In the **Distribution** list, select the distribution for which you want to display a list of top referrers.
5. Click **Update**.

Topics

- [Downloading Data in CSV Format \(p. 415\)](#)
- [How Data in the Top Referrers Report Is Related to Data in the CloudFront Access Logs \(p. 416\)](#)

Downloading Data in CSV Format

You can download the Top Referrers report in CSV format. This section explains how to download the report and describes the values in the report.

To download the Top Referrers report in CSV format

1. While viewing the Top Referrers report, click **CSV**.
2. In the **Opening *file name*** dialog box, choose whether to open or save the file.

Information About the Report

The first few rows of the report include the following information:

Version

The version of the format for this CSV file.

Report

The name of the report.

DistributionID

The ID of the distribution that you ran the report for, or `ALL` if you ran the report for all distributions.

StartDateUTC

The beginning of the date range for which you ran the report, in Coordinated Universal Time (UTC).

EndDateUTC

The end of the date range for which you ran the report, in Coordinated Universal Time (UTC).

GeneratedTimeUTC

The date and time on which you ran the report, in Coordinated Universal Time (UTC).

Data in the Top Referrers Report

The report includes the following values:

DistributionID

The ID of the distribution that you ran the report for, or `ALL` if you ran the report for all distributions.

FriendlyName

An alternate domain name (CNAME) for the distribution, if any. If a distribution has no alternate domain names, the list includes an origin domain name for the distribution.

Referrer

The domain name of the referrer.

The total number of requests from the domain name in the `Referrer` column.

RequestsPct

The number of requests submitted by the referrer as a percentage of the total number of requests during the specified period.

How Data in the Top Referrers Report Is Related to Data in the CloudFront Access Logs

The following list shows how values in the Top Referrers report in the CloudFront console correspond with values in CloudFront access logs. For more information about CloudFront access logs, see [Configuring and Using Standard Logs \(Access Logs\)](#) (p. 439).

Referrer

The domain name of the referrer. In access logs, referrers are listed in the `cs(Referrer)` column.

Request Count

The total number of requests from the domain name in the **Referrer** column. This value generally corresponds closely with the number of `GET` requests from the referrer in CloudFront access logs.

Request %

The number of requests submitted by the referrer as a percentage of the total number of requests during the specified period. If you have more than 25 referrers, then you can't calculate **Request %** based on the data in this table because the **Request Count** column doesn't include all of the requests during the specified period.

CloudFront Usage Reports

The Amazon CloudFront console can display a graphical representation of your CloudFront usage that is based on a subset of the usage report data. You can display charts for a specified date range in the last 60 days, with data points every hour or every day. You can usually view data about requests that CloudFront received as recently as four hours ago, but data can occasionally be delayed by as much as 24 hours.

For more information, see [How the Usage Charts Are Related to Data in the CloudFront Usage Report](#) (p. 419).

To display CloudFront usage charts

1. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
2. In **navigation** pane, click **Usage Reports**.

3. In the **CloudFront Usage Reports** pane, for **Start Date** and **End Date**, select the date range for which you want to display usage charts. Available ranges depend on the value that you select for **Granularity**:
 - **Daily** — To display charts with one data point per day, select any date range in the previous 60 days.
 - **Hourly** — To display charts with one data point every hour, select any date range of up to 14 days within the previous 60 days.

Dates and times are in Coordinated Universal Time (UTC).

4. For **Granularity**, specify whether to display one data point per day or one data point per hour in the charts. If you specify a date range greater than 14 days, the option to specify one data point per hour is not available.
5. For **Billing Region**, choose the CloudFront billing region that has the data you want to view, or choose **All Regions**. Usage charts include data for requests that CloudFront processes in edge locations in the specified region. The region where CloudFront processes requests might or might not correspond with the location of your users.

Select only regions that are included in the price class for your distribution; otherwise, the usage charts probably won't contain any data. For example, if you chose Price Class 200 for your distribution, the South America and Australia billing regions are not included, so CloudFront generally won't process your requests from those regions. For more information about price classes, see [Choosing the price class for a CloudFront distribution \(p. 10\)](#).

6. In the **Distribution** list, select the distributions for which you want to display data in the usage charts:
 - **An individual web distribution** — The charts display data for the selected CloudFront distribution. The **Distribution** list displays the distribution ID and alternate domain names (CNAMEs) for the distribution, if any. If a distribution has no alternate domain names, the list includes origin domain names for the distribution.
 - **All Web Distributions (excludes deleted)** — The charts display summed data for all web distributions that are associated with the current AWS account, excluding web distributions that you have deleted.
 - **All Deleted Distributions** — The charts display summed data for all web distributions that are associated with the current AWS account and that were deleted in the last 60 days.
7. Click **Update Graphs**.
8. To view data for a daily or hourly data point within a chart, move your mouse pointer over the data point.
9. For charts that show data transferred, note that you can change the vertical scale to gigabytes, megabytes, or kilobytes for each chart.

Topics

- [Downloading Data in CSV Format \(p. 417\)](#)
- [How the Usage Charts Are Related to Data in the CloudFront Usage Report \(p. 419\)](#)

Downloading Data in CSV Format

You can download the Usage report in CSV format. This section explains how to download the report and describes the values in the report.

To download the Usage report in CSV format

1. While viewing the Usage report, click **CSV**.

2. In the **Opening *file name*** dialog box, choose whether to open or save the file.

Information About the Report

The first few rows of the report include the following information:

Version

The version of the format for this CSV file.

Report

The name of the report.

DistributionID

The ID of the distribution that you ran the report for, `ALL` if you ran the report for all distributions, or `ALL_DELETED` if you ran the report for all deleted distributions.

StartDateUTC

The beginning of the date range for which you ran the report, in Coordinated Universal Time (UTC).

EndDateUTC

The end of the date range for which you ran the report, in Coordinated Universal Time (UTC).

GeneratedTimeUTC

The date and time on which you ran the report, in Coordinated Universal Time (UTC).

Granularity

Whether each row in the report represents one hour or one day.

BillingRegion

The continent that viewer requests originated from, or `ALL`, if you chose to download the report for all billing regions.

Data in the Usage Report

The report includes the following values:

DistributionID

The ID of the distribution that you ran the report for, `ALL` if you ran the report for all distributions, or `ALL_DELETED` if you ran the report for all deleted distributions.

FriendlyName

An alternate domain name (CNAME) for the distribution, if any. If a distribution has no alternate domain names, the list includes an origin domain name for the distribution.

BillingRegion

The CloudFront billing region that you ran the report for, or `ALL`.

TimeBucket

The hour or the day that data applies to, in Coordinated Universal Time (UTC).

HTTP

The number of HTTP requests that CloudFront responded to from edge locations in the selected region during each time interval for the specified CloudFront distribution. Values include:

- The number of `GET` and `HEAD` requests, which cause CloudFront to transfer data to your users
- The number of `DELETE`, `OPTIONS`, `PATCH`, `POST`, and `PUT` requests, which cause CloudFront to transfer data to your origin

HTTPS

The number of HTTPS requests that CloudFront responded to from edge locations in the selected region during each time interval for the specified CloudFront distribution. Values include:

- The number of `GET` and `HEAD` requests, which cause CloudFront to transfer data to your users
- The number of `DELETE`, `OPTIONS`, `PATCH`, `POST`, and `PUT` requests, which cause CloudFront to transfer data to your origin

HTTPBytes

The total amount of data transferred over HTTP from CloudFront edge locations in the selected billing region during the time period for the specified CloudFront distribution. Values include:

- Data transferred from CloudFront to your users in response to `GET` and `HEAD` requests
- Data transferred from CloudFront to your origin for `DELETE`, `OPTIONS`, `PATCH`, `POST`, and `PUT` requests
- Data transferred from CloudFront to your users in response to `DELETE`, `OPTIONS`, `PATCH`, `POST`, and `PUT` requests

HTTPSBytes

The total amount of data transferred over HTTPS from CloudFront edge locations in the selected billing region during the time period for the specified CloudFront distribution. Values include:

- Data transferred from CloudFront to your users in response to `GET` and `HEAD` requests
- Data transferred from CloudFront to your origin for `DELETE`, `OPTIONS`, `PATCH`, `POST`, and `PUT` requests
- Data transferred from CloudFront to your users in response to `DELETE`, `OPTIONS`, `PATCH`, `POST`, and `PUT` requests

BytesIn

The total amount of data transferred from CloudFront to your origin for `DELETE`, `OPTIONS`, `PATCH`, `POST`, and `PUT` requests in the selected region during each time interval for the specified CloudFront distribution.

BytesOut

The total amount of data transferred over HTTP and HTTPS from CloudFront to your users in the selected region during each time interval for the specified CloudFront distribution. Values include:

- Data transferred from CloudFront to your users in response to `GET` and `HEAD` requests
- Data transferred from CloudFront to your users in response to `DELETE`, `OPTIONS`, `PATCH`, `POST`, and `PUT` requests

How the Usage Charts Are Related to Data in the CloudFront Usage Report

The following list shows how the usage charts in the CloudFront console correspond with values in the **Usage Type** column in the CloudFront usage report.

Topics

- [Number of Requests \(p. 420\)](#)
- [Data Transferred by Protocol \(p. 420\)](#)
- [Data Transferred by Destination \(p. 420\)](#)

Number of Requests

This chart shows the total number of requests that CloudFront responds to from edge locations in the selected region during each time interval for the specified CloudFront distribution, separated by protocol (HTTP or HTTPS) and type (static, dynamic, or proxy).

Number of HTTP Requests

- **region-Requests-HTTP-Static:** Number of HTTP GET and HEAD requests served for objects with TTL \geq 3600 seconds
- **region-Requests-HTTP-Dynamic:** Number of HTTP GET and HEAD requests served for objects with TTL $<$ 3600 seconds
- **region-Requests-HTTP-Proxy:** Number of HTTP DELETE, OPTIONS, PATCH, POST, and PUT requests that CloudFront forwards to your origin

Number of HTTPS Requests

- **region-Requests-HTTPS-Static:** Number of HTTPS GET and HEAD requests served for objects with TTL \geq 3600 seconds
- **region-Requests-HTTPS-Dynamic:** Number of HTTPS GET and HEAD requests served for objects with TTL $<$ 3600 seconds
- **region-Requests-HTTPS-Proxy:** Number of HTTPS DELETE, OPTIONS, PATCH, POST, and PUT requests that CloudFront forwards to your origin

Data Transferred by Protocol

This chart shows the total amount of data transferred from CloudFront edge locations in the selected region during each time interval for the specified CloudFront distribution, separated by protocol (HTTP or HTTPS), type (static, dynamic, or proxy), and destination (users or origin).

Data Transferred over HTTP

- **region-Out-Bytes-HTTP-Static:** Bytes served via HTTP for objects with TTL \geq 3600 seconds
- **region-Out-Bytes-HTTP-Dynamic:** Bytes served via HTTP for objects with TTL $<$ 3600 seconds
- **region-Out-Bytes-HTTP-Proxy:** Bytes returned from CloudFront to viewers via HTTP in response to DELETE, OPTIONS, PATCH, POST, and PUT requests
- **region-Out-OBytes-HTTP-Proxy:** Total bytes transferred via HTTP from CloudFront edge locations to your origin in response to DELETE, OPTIONS, PATCH, POST, and PUT requests

Data Transferred over HTTPS

- **region-Out-Bytes-HTTPS-Static:** Bytes served via HTTPS for objects with TTL \geq 3600 seconds
- **region-Out-Bytes-HTTPS-Dynamic:** Bytes served via HTTPS for objects with TTL $<$ 3600 seconds
- **region-Out-Bytes-HTTPS-Proxy:** Bytes returned from CloudFront to viewers via HTTPS in response to DELETE, OPTIONS, PATCH, POST, and PUT requests
- **region-Out-OBytes-HTTPS-Proxy:** Total bytes transferred via HTTPS from CloudFront edge locations to your origin in response to DELETE, OPTIONS, PATCH, POST, and PUT requests

Data Transferred by Destination

This chart shows the total amount of data transferred from CloudFront edge locations in the selected region during each time interval for the specified CloudFront distribution, separated by destination (users or origin), protocol (HTTP or HTTPS), and type (static, dynamic, or proxy).

Data Transferred from CloudFront to Your Users

- **region-Out-Bytes-HTTP-Static:** Bytes served via HTTP for objects with TTL \geq 3600 seconds
- **region-Out-Bytes-HTTPS-Static:** Bytes served via HTTPS for objects with TTL \geq 3600 seconds

- **region-Out-Bytes-HTTP-Dynamic:** Bytes served via HTTP for objects with TTL < 3600 seconds
- **region-Out-Bytes-HTTPS-Dynamic:** Bytes served via HTTPS for objects with TTL < 3600 seconds
- **region-Out-Bytes-HTTP-Proxy:** Bytes returned from CloudFront to viewers via HTTP in response to DELETE, OPTIONS, PATCH, POST, and PUT requests
- **region-Out-Bytes-HTTPS-Proxy:** Bytes returned from CloudFront to viewers via HTTPS in response to DELETE, OPTIONS, PATCH, POST, and PUT requests

Data Transferred from CloudFront to Your Origin

- **region-Out-OBytes-HTTP-Proxy:** Total bytes transferred via HTTP from CloudFront edge locations to your origin in response to DELETE, OPTIONS, PATCH, POST, and PUT requests
- **region-Out-OBytes-HTTPS-Proxy:** Total bytes transferred via HTTPS from CloudFront edge locations to your origin in response to DELETE, OPTIONS, PATCH, POST, and PUT requests

CloudFront Viewers Reports

The CloudFront console can display four reports about the physical devices (desktop computers, mobile devices) and about the viewers (typically web browsers) that are accessing your content:

- **Devices** – The type of the devices that your users use most frequently to access your content, for example, Desktop or Mobile.
- **Browsers** – The name (or the name and version) of the browsers that your users use most frequently to access your content, for example, Chrome or Firefox. The report lists the top 10 browsers.
- **Operating Systems** – The name (or the name and version) of the operating system that viewers run on most frequently when accessing your content, for example, Linux, Mac OS X, or Windows. The report lists the top 10 operating systems.
- **Locations** – The locations, by country or by U.S. state/territory, of the viewers that access your content most frequently. The report lists the top 50 countries or U.S. states/territories.

You can display all four Viewers reports for any date range in the previous 60 days. For the Locations report, you can also display the report with data points every hour for any date range of up to 14 days in the previous 60 days.

Note

You don't need to enable access logging to view Viewers charts and reports.

Topics

- [Displaying Viewers Charts and Reports \(p. 421\)](#)
- [Downloading Data in CSV Format \(p. 422\)](#)
- [How Data in the Locations Report Is Related to Data in the CloudFront Access Logs \(p. 427\)](#)

Displaying Viewers Charts and Reports

To display CloudFront Viewers charts and reports, perform the following procedure.

To display CloudFront Viewers charts and reports

1. Sign in to the AWS Management Console and open the CloudFront console at <https://console.aws.amazon.com/cloudfront/>.
2. In the navigation pane, click **Viewers**.
3. In the **CloudFront Viewers** pane, for **Start Date** and **End Date**, select the date range for which you want to display viewer charts and reports.

For the Locations chart, available ranges depend on the value that you select for **Granularity**:

- **Daily** – To display charts with one data point per day, select any date range in the previous 60 days.
- **Hourly** – To display charts with one data point every hour, select any date range of up to 14 days within the previous 60 days.

Dates and times are in Coordinated Universal Time (UTC).

4. (Browsers and Operating Systems charts only) For **Grouping**, specify whether you want to group browsers and operating systems by name (Chrome, Firefox) or by name and version (Chrome 40.0, Firefox 35.0).
5. (Locations chart only) For **Granularity**, specify whether to display one data point per day or one data point per hour in the charts. If you specify a date range greater than 14 days, the option to specify one data point per hour is not available.
6. (Locations chart only) For **Details**, specify whether to display the top locations by countries or by U.S. states.
7. In the **Distribution** list, select the distribution for which you want to display data in the usage charts:
 - **An individual web distribution** – The charts display data for the selected CloudFront web distribution. The **Distribution** list displays the distribution ID and an alternate domain name (CNAME) for the distribution, if any. If a distribution has no alternate domain names, the list includes an origin domain name for the distribution.
 - **All Web Distributions (excludes deleted)** – The charts display summed data for all web distributions that are associated with the current AWS account, excluding web distributions that you have deleted.
8. Click **Update**.
9. To view data for a daily or hourly data point within a chart, move your mouse pointer over the data point.

Downloading Data in CSV Format

You can download each of the Viewer reports in CSV format. This section explains how to download the reports and describes the values in the report.

To download the Viewer reports in CSV format

1. While viewing the Viewer report, click **CSV**.
2. Choose the data that you want to download, for example, **Devices** or **Devices Trends**.
3. In the **Opening file name** dialog box, choose whether to open or save the file.

Topics

- [Information About the Reports \(p. 423\)](#)
- [Devices Report \(p. 423\)](#)
- [Device Trends Report \(p. 424\)](#)
- [Browsers Report \(p. 424\)](#)
- [Browser Trends Report \(p. 425\)](#)
- [Operating Systems Report \(p. 425\)](#)
- [Operating System Trends Report \(p. 426\)](#)
- [Locations Report \(p. 426\)](#)

- [Location Trends Report \(p. 427\)](#)

Information About the Reports

The first few rows of each report includes the following information:

Version

The version of the format for this CSV file.

Report

The name of the report.

DistributionID

The ID of the distribution that you ran the report for, or `ALL` if you ran the report for all web distributions.

StartDateUTC

The beginning of the date range for which you ran the report, in Coordinated Universal Time (UTC).

EndDateUTC

The end of the date range for which you ran the report, in Coordinated Universal Time (UTC).

GeneratedTimeUTC

The date and time on which you ran the report, in Coordinated Universal Time (UTC).

Grouping (Browsers and Operating Systems Reports Only)

Whether the data is grouped by the name or by the name and version of the browser or operating system.

Granularity

Whether each row in the report represents one hour or one day.

Details (Locations Report Only)

Whether requests are listed by country or by U.S. state.

Devices Report

The report includes the following values:

DistributionID

The ID of the distribution that you ran the report for, or `ALL` if you ran the report for all distributions.

FriendlyName

An alternate domain name (CNAME) for the distribution, if any. If a distribution has no alternate domain names, the list includes an origin domain name for the distribution.

Requests

The number of requests that CloudFront received from each type of device.

RequestsPct

The number of requests that CloudFront received from each type of device as a percentage of the total number of requests that CloudFront received from all devices.

Device Trends Report

The report includes the following values:

DistributionID

The ID of the distribution that you ran the report for, or `ALL` if you ran the report for all distributions.

FriendlyName

An alternate domain name (CNAME) for the distribution, if any. If a distribution has no alternate domain names, the list includes an origin domain name for the distribution.

TimeBucket

The hour or the day that the data applies to, in Coordinated Universal Time (UTC).

Desktop

The number of requests that CloudFront received from desktop computers during the period.

Mobile

The number of requests that CloudFront received from mobile devices during the period. Mobile devices can include both tablets and mobile phones. If CloudFront can't determine whether a request originated from a mobile device or a tablet, it's counted in the `Mobile` column.

Smart-TV

The number of requests that CloudFront received from smart TVs during the period.

Tablet

The number of requests that CloudFront received from tablets during the period. If CloudFront can't determine whether a request originated from a mobile device or a tablet, it's counted in the `Mobile` column.

Unknown

Requests for which the value of the `User-Agent` HTTP header was not associated with one of the standard device types, for example, `Desktop` or `Mobile`.

Empty

The number of requests that CloudFront received that didn't include a value in the HTTP `User-Agent` header during the period.

Browsers Report

The report includes the following values:

DistributionID

The ID of the distribution that you ran the report for, or `ALL` if you ran the report for all distributions.

FriendlyName

An alternate domain name (CNAME) for the distribution, if any. If a distribution has no alternate domain names, the list includes an origin domain name for the distribution.

Group

The browser or the browser and version that CloudFront received requests from, depending on the value of `Grouping`. In addition to browser names, possible values include the following:

- **Bot/Crawler** – primarily requests from search engines that are indexing your content.
- **Empty** – requests for which the value of the `User-Agent` HTTP header was empty.
- **Other** – browsers that CloudFront identified but that aren't among the most popular. If `Bot/Crawler`, `Empty`, and/or `Unknown` don't appear among the first nine values, then they're also included in `Other`.
- **Unknown** – requests for which the value of the `User-Agent` HTTP header was not associated with a standard browser. Most requests in this category come from custom applications or scripts.

Requests

The number of requests that CloudFront received from each type of browser.

RequestsPct

The number of requests that CloudFront received from each type of browser as a percentage of the total number of requests that CloudFront received during the time period.

Browser Trends Report

The report includes the following values:

DistributionID

The ID of the distribution that you ran the report for, or `ALL` if you ran the report for all distributions.

FriendlyName

An alternate domain name (CNAME) for the distribution, if any. If a distribution has no alternate domain names, the list includes an origin domain name for the distribution.

TimeBucket

The hour or the day that the data applies to, in Coordinated Universal Time (UTC).

(Browsers)

The remaining columns in the report list the browsers or the browsers and their versions, depending on the value of `Grouping`. In addition to browser names, possible values include the following:

- **Bot/Crawler** – primarily requests from search engines that are indexing your content.
- **Empty** – requests for which the value of the `User-Agent` HTTP header was empty.
- **Other** – browsers that CloudFront identified but that aren't among the most popular. If `Bot/Crawler`, `Empty`, and/or `Unknown` don't appear among the first nine values, then they're also included in `Other`.
- **Unknown** – requests for which the value of the `User-Agent` HTTP header was not associated with a standard browser. Most requests in this category come from custom applications or scripts.

Operating Systems Report

The report includes the following values:

DistributionID

The ID of the distribution that you ran the report for, or `ALL` if you ran the report for all distributions.

FriendlyName

An alternate domain name (CNAME) for the distribution, if any. If a distribution has no alternate domain names, the list includes an origin domain name for the distribution.

Group

The operating system or the operating system and version that CloudFront received requests from, depending on the value of `Grouping`. In addition to operating system names, possible values include the following:

- **Bot/Crawler** – primarily requests from search engines that are indexing your content.
- **Empty** – requests for which the value of the `User-Agent` HTTP header was empty.
- **Other** – operating systems that CloudFront identified but that aren't among the most popular. If `Bot/Crawler`, `Empty`, and/or `Unknown` don't appear among the first nine values, then they're also included in `Other`.
- **Unknown** – requests for which the value of the `User-Agent` HTTP header was not associated with a standard browser. Most requests in this category come from custom applications or scripts.

Requests

The number of requests that CloudFront received from each type of operating system.

RequestsPct

The number of requests that CloudFront received from each type of operating system as a percentage of the total number of requests that CloudFront received during the time period.

Operating System Trends Report

The report includes the following values:

DistributionID

The ID of the distribution that you ran the report for, or `ALL` if you ran the report for all distributions.

FriendlyName

An alternate domain name (CNAME) for the distribution, if any. If a distribution has no alternate domain names, the list includes an origin domain name for the distribution.

TimeBucket

The hour or the day that the data applies to, in Coordinated Universal Time (UTC).

(Operating systems)

The remaining columns in the report list the operating systems or the operating systems and their versions, depending on the value of `Grouping`. In addition to operating system names, possible values include the following:

- **Bot/Crawler** – primarily requests from search engines that are indexing your content.
- **Empty** – requests for which the value of the `User-Agent` HTTP header was empty.
- **Other** – operating systems that CloudFront identified but that aren't among the most popular. If `Bot/Crawler`, `Empty`, and/or `Unknown` don't appear among the first nine values, then they're also included in `Other`.
- **Unknown** – requests for which the operating system isn't specified in the `User-Agent` HTTP header.

Locations Report

The report includes the following values:

DistributionID

The ID of the distribution that you ran the report for, or `ALL` if you ran the report for all distributions.

FriendlyName

An alternate domain name (CNAME) for the distribution, if any. If a distribution has no alternate domain names, the list includes an origin domain name for the distribution.

LocationCode

The abbreviation for the location that CloudFront received requests from. For more information about possible values, see the description of Location in [How Data in the Locations Report Is Related to Data in the CloudFront Access Logs](#) (p. 427).

LocationName

The name of the location that CloudFront received requests from.

Requests

The number of requests that CloudFront received from each location.

RequestsPct

The number of requests that CloudFront received from each location as a percentage of the total number of requests that CloudFront received from all locations during the time period.

TotalBytes

The number of bytes that CloudFront served to viewers in this country or state, for the specified distribution and period.

Location Trends Report

The report includes the following values:

DistributionID

The ID of the distribution that you ran the report for, or `ALL` if you ran the report for all distributions.

FriendlyName

An alternate domain name (CNAME) for the distribution, if any. If a distribution has no alternate domain names, the list includes an origin domain name for the distribution.

TimeBucket

The hour or the day that the data applies to, in Coordinated Universal Time (UTC).

(Locations)

The remaining columns in the report list the locations that CloudFront received requests from. For more information about possible values, see the description of Location in [How Data in the Locations Report Is Related to Data in the CloudFront Access Logs](#) (p. 427).

How Data in the Locations Report Is Related to Data in the CloudFront Access Logs

The following list shows how data in the Locations report in the CloudFront console corresponds with values in CloudFront access logs. For more information about CloudFront access logs, see [Configuring and Using Standard Logs \(Access Logs\)](#) (p. 439).

Location

The country or U.S. state that the viewer is in. In access logs, the `c-ip` column contains the IP address of the device that the viewer is running on. We use geolocation data to identify the geographic location of the device based on the IP address.

If you're displaying the **Locations** report by country, note that the country list is based on [ISO 3166-2, Codes for the representation of names of countries and their subdivisions – Part 2: Country subdivision code](#). The country list includes the following additional values:

- **Anonymous Proxy** – The request originated from an anonymous proxy.
- **Satellite Provider** – The request originated from a satellite provider that provides internet service to multiple countries. Users might be in countries with a high risk of fraud.
- **Europe (Unknown)** – The request originated from an IP in a block that is used by multiple European countries. The country that the request originated from cannot be determined. CloudFront uses **Europe (Unknown)** as the default.
- **Asia/Pacific (Unknown)** – The request originated from an IP in a block that is used by multiple countries in the Asia/Pacific region. The country that the request originated from cannot be determined. CloudFront uses **Asia/Pacific (Unknown)** as the default.

If you display the **Locations** report by U.S. state, note that the report can include U.S. territories and U.S. Armed Forces regions.

Note

If CloudFront can't determine a user's location, the location will appear as Unknown in viewer reports.

Request Count

The total number of requests from the country or U.S. state that the viewer is in, for the specified distribution and period. This value generally corresponds closely with the number of `GET` requests from IP addresses in that country or state in CloudFront access logs.

Request %

One of the following, depending on the value that you selected for **Details**:

- **Countries** – The requests from this country as a percentage of the total number of requests.
- **U.S. States** – The requests from this state as a percentage of the total number of requests from the United States.

If requests came from more than 50 countries, then you can't calculate **Request %** based on the data in this table because the **Request Count** column doesn't include all of the requests during the specified period.

Bytes

The number of bytes that CloudFront served to viewers in this country or state, for the specified distribution and period. To change the display of data in this column to KB, MB, or GB, click the link in the column heading.

Monitoring CloudFront with Amazon CloudWatch

Amazon CloudFront is integrated with Amazon CloudWatch, and automatically publishes six operational metrics per distribution, which are displayed in a set of graphs in the [CloudFront console](#), or accessible by using the CloudFront API or CLI. Each Lambda@Edge function associated with a distribution also publishes operational metrics, some of which you can view in the CloudFront console. The CloudFront metrics don't count against [CloudWatch quotas \(formerly known as limits\)](#) and don't incur any additional cost.

In addition to the default metrics, you can enable additional metrics for an additional cost. The additional metrics apply to CloudFront distributions, and must be enabled for each distribution separately. For more information about the cost, see [Estimating cost for the additional CloudFront metrics](#) (p. 432).

Viewing these metrics can help you troubleshoot, track, and debug issues. To view these metrics in the CloudFront console, see the [Monitoring page](#). To view graphs about the activity for a specific CloudFront distribution or Lambda@Edge function, choose one, and then choose to view the metrics.

You can also set alarms based on these metrics in the CloudFront console, or in the CloudWatch console, API, or CLI ([standard CloudWatch pricing](#) applies). For example, you can set an alarm based on the `5xxErrorRate` metric, which represents the percentage of all viewer requests for which the response's HTTP status code is in the range of 500 to 599, inclusive. When the error rate reaches a certain value for a certain amount of time—for example, 5% of requests for 5 continuous minutes—the alarm is triggered. You specify the alarm's value and its time unit when you create the alarm. For more information, see [Setting alarms](#) (p. 432).

Topics

- [Viewing CloudFront and Lambda@Edge metrics](#) (p. 429)
- [Setting alarms to receive notifications](#) (p. 432)
- [Downloading data in CSV format](#) (p. 433)
- [Getting metrics using the CloudWatch API](#) (p. 435)

Viewing CloudFront and Lambda@Edge metrics

You can view operational metrics about your CloudFront distributions and Lambda@Edge functions in the CloudFront console. To view these metrics, see the [Monitoring page in the CloudFront console](#). To view graphs about the activity for a specific CloudFront distribution or Lambda@Edge function, choose one, and then choose to view the metrics.

Topics

- [Viewing the default CloudFront distribution metrics](#) (p. 429)
- [Viewing additional CloudFront distribution metrics](#) (p. 430)
- [Viewing the default Lambda@Edge function metrics](#) (p. 432)

Viewing the default CloudFront distribution metrics

The following default metrics are included for all CloudFront distributions, at no additional cost:

Requests

The total number of viewer requests received by CloudFront, for all HTTP methods and for both HTTP and HTTPS requests.

Bytes downloaded

The total number of bytes downloaded by viewers for `GET`, `HEAD`, and `OPTIONS` requests.

Bytes uploaded

The total number of bytes that viewers uploaded to your origin with CloudFront, using `POST` and `PUT` requests.

4xx error rate

The percentage of all viewer requests for which the response's HTTP status code is `4xx`.

5xx error rate

The percentage of all viewer requests for which the response's HTTP status code is 5xx.

Total error rate

The percentage of all viewer requests for which the response's HTTP status code is 4xx or 5xx.

These metrics are shown in graphs for each CloudFront distribution on the [Monitoring page in the CloudFront console](#). On each graph, the totals are displayed at 1-minute granularity. In addition to viewing the graphs, you can also [download metrics reports as CSV files \(p. 433\)](#).

You can customize the graphs by doing the following:

- To change the time range for the information displayed in the graphs, choose 1h (1 hour), 3h (3 hours), or another range, or specify a custom range.
- To change how often CloudFront updates the information in the graph, choose the down arrow next to the refresh icon, and then choose a refresh rate. The default refresh rate is 1 minute, but you can choose 10 seconds, 2 minutes, or other options.

To view CloudFront graphs in the CloudWatch console, choose **Add to dashboard**.

Viewing additional CloudFront distribution metrics

In addition to the default metrics, you can enable additional metrics for an additional cost. For more information about the cost, see [Estimating cost for the additional CloudFront metrics \(p. 432\)](#).

These additional metrics must be enabled for each distribution separately:

Cache hit rate

The percentage of all cacheable requests for which CloudFront served the content from its cache. HTTP POST and PUT requests, and errors, are not considered cacheable requests.

Origin latency

The total time spent from when CloudFront receives a request to when it starts providing a response to the network (not the viewer), for requests that are served from the origin, not the CloudFront cache. This is also known as *first byte latency*, or *time-to-first-byte*.

Error rate by status code

The percentage of all viewer requests for which the response's HTTP status code is a particular code in the 4xx or 5xx range. This metric is available for all of the following error codes: 401, 403, 404, 502, 503, and 504.

Enabling additional metrics

You can enable additional metrics in the CloudFront console, with the AWS Command Line Interface (AWS CLI), or with the CloudFront API.

Enabling additional metrics (console)

To enable additional metrics

1. Sign in to the AWS Management Console and open the [Monitoring page in the CloudFront console](#).
2. Choose the distribution to enable additional metrics for, and then choose **View distribution metrics**.
3. Choose **Enable additional metrics**.

4. In the **Enable additional metrics** window, choose **Enable**, and then choose **Save**.

After you enable the additional metrics, they are shown in graphs. On each graph, the totals are displayed at 1-minute granularity. In addition to viewing the graphs, you can also [download metrics reports as CSV files](#) (p. 433).

You can customize the graphs by doing the following:

- To change the time range for the information displayed in the graphs, choose 1h (1 hour), 3h (3 hours), or another range, or specify a custom range.
- To change how often CloudFront updates the information in the graph, choose the down arrow next to the refresh icon, and then choose a refresh rate. The default refresh rate is 1 minute, but you can choose 10 seconds, 2 minutes, or other options.

To view CloudFront graphs in the CloudWatch console, choose **Add to dashboard**.

Enabling additional metrics (AWS CLI)

To manage additional metrics with the AWS Command Line Interface (AWS CLI), use one of the following commands.

To enable additional metrics for a distribution

- Use the **create-monitoring-subscription** command, as in the following example. Replace **EDFDVBD6EXAMPLE** with the ID of the distribution that you are enabling additional metrics for.

```
aws cloudfront create-monitoring-subscription --  
distribution-id EDFDVBD6EXAMPLE --monitoring-subscription  
RealtimeMetricsSubscriptionConfig={RealtimeMetricsSubscriptionStatus=Enabled}
```

To see whether additional metrics are enabled for a distribution

- Use the **get-monitoring-subscription** command, as in the following example. Replace **EDFDVBD6EXAMPLE** with the ID of the distribution that you are checking.

```
aws cloudfront get-monitoring-subscription --distribution-id EDFDVBD6EXAMPLE
```

To disable additional metrics for a distribution

- Use the **delete-monitoring-subscription** command, as in the following example. Replace **EDFDVBD6EXAMPLE** with the ID of the distribution that you are disabling additional metrics for.

```
aws cloudfront delete-monitoring-subscription --distribution-id EDFDVBD6EXAMPLE
```

Enabling additional metrics (API)

To manage additional metrics with the CloudFront API, use one of the following API operations.

- To enable additional metrics for a distribution, use [CreateMonitoringSubscription](#).
- To see whether additional metrics are enabled for a distribution, use [GetMonitoringSubscription](#).

- To disable additional metrics for a distribution, use [DeleteMonitoringSubscription](#).

For more information about these API calls, see the API reference documentation for your AWS SDK or other API client.

Estimating cost for the additional CloudFront metrics

When you enable additional metrics for a distribution, CloudFront sends up to 8 metrics to CloudWatch in the US East (N. Virginia) Region. CloudWatch charges a low, fixed rate for each metric. This rate is charged only once per month, per metric (up to 8 metrics per distribution). This is a fixed rate, so your cost remains the same regardless of the number of requests or responses that the CloudFront distribution receives or sends. For the per-metric rate, see the [Amazon CloudWatch pricing page](#) and the [CloudWatch pricing calculator](#). Additional API charges apply when you retrieve the metrics with the CloudWatch API.

Viewing the default Lambda@Edge function metrics

The following default metrics are shown in graphs for each Lambda@Edge function on the [Monitoring page in the CloudFront console](#):

- 5xx error rate for Lambda@Edge
- Lambda execution errors
- Lambda invalid responses
- Lambda throttles

The graphs include the number of invocations, errors, throttles, and so on. On each graph, the totals are displayed at 1-minute granularity, grouped by AWS Region.

If you see a spike in errors that you want to investigate, you can choose a function and then view log files by AWS Region, until you determine which function is causing the problems and in which AWS Region. For more information about troubleshooting Lambda@Edge errors, see:

- [How to Determine the Type of Failure \(p. 345\)](#)
- [Four Steps for Debugging your Content Delivery on AWS](#)

You can customize the graphs by doing the following:

- To change the time range for the information displayed in the graphs, choose 1h (1 hour), 3h (3 hours), or another range, or specify a custom range.
- To change how often CloudFront updates the information in the graph, choose the down arrow next to the refresh icon, and then choose a refresh rate. The default refresh rate is 1 minute, but you can choose 10 seconds, 2 minutes, or other options.

To view the graphs in the CloudWatch console, choose **Add to dashboard**. You must use the US East (N. Virginia) Region (us-east-1) to view the graphs in the CloudWatch console.

Setting alarms to receive notifications

In the CloudFront console, you can set alarms to notify you by Amazon Simple Notification Service (Amazon SNS) based on specific CloudFront metrics. You can set an alarm on the [Alarms page in the CloudFront console](#).

When you create an alarm in the console, you specify the following values:

Metric

The metric for which you are creating the alarm.

Distribution

The CloudFront distribution for which you are creating the alarm.

Name of alarm

A name for the alarm.

Send a notification to

The Amazon SNS topic to send notification to if this metric triggers an alarm.

Whenever *<metric>* *<operator>* *<value>*

Specify when CloudWatch should trigger an alarm and send a notification to the Amazon SNS topic. For example, to receive a notification when the 5xx error rate exceeds 1%, specify the following:

Whenever Average of 5xxErrorRate > 1

Note the following about specifying values:

- Enter only whole numbers without punctuation. For example, to specify one thousand, enter **1000**.
- For 4xx, 5xx, and total error rates, the value that you specify is a percentage.
- For requests, bytes downloaded, and bytes uploaded, the value that you specify is units. For example, 1073742000 bytes.

For at least *<number>* **consecutive periods of** *<time period>*

Specify how many consecutive time periods of the specified duration the metric must meet the criteria before CloudWatch triggers an alarm. When you choose a value, aim for an appropriate balance between a value that does not alarm for temporary or fleeting problems, but does alarm for sustained or real problems.

Downloading data in CSV format

You can download the CloudWatch metrics data for a CloudFront distribution in CSV format. You can download the data when you **View distribution metrics** for a particular distribution in the [CloudFront console](#).

Information about the report

The first few rows of the report include the following information:

Version

The CloudFront reporting version.

Report

The name of the report.

DistributionID

The ID of the distribution for which you ran the report.

StartDateUTC

The beginning of the date range for which you ran the report, in Coordinated Universal Time (UTC).

EndDateUTC

The end of the date range for which you ran the report, in Coordinated Universal Time (UTC).

GeneratedTimeUTC

The date and time on which you ran the report, in Coordinated Universal Time (UTC).

Granularity

The time period for each row in the report, for example, `ONE_MINUTE`.

Data in the metrics report

The report includes the following values:

DistributionID

The ID of the distribution for which you ran the report.

FriendlyName

An alternate domain name (CNAME) for the distribution, if any. If a distribution has no alternate domain names, the list includes an origin domain name for the distribution.

TimeBucket

The hour or the day that the data applies to, in Coordinated Universal Time (UTC).

Requests

The total number of requests for all HTTP status codes (for example, 200, 404, and so on) and all methods (for example, GET, HEAD, POST, and so on) during the time period.

BytesDownloaded

The number of bytes that viewers downloaded for the specified distribution during the time period.

BytesUploaded

The number of bytes that viewers uploaded for the specified distribution during the time period.

TotalErrorRatePct

The percentage of requests for which the HTTP status code was a 4xx or 5xx error for the specified distribution during the time period.

4xxErrorRatePct

The percentage of requests for which the HTTP status code was a 4xx error for the specified distribution during the time period.

5xxErrorRatePct

The percentage of requests for which the HTTP status code was a 5xx error for the specified distribution during the time period.

If you have [enabled additional metrics \(p. 430\)](#) for your distribution, then the report also includes the following additional values:

401ErrorRatePct

The percentage of requests for which the HTTP status code was a 401 error for the specified distribution during the time period.

403ErrorRatePct

The percentage of requests for which the HTTP status code was a 403 error for the specified distribution during the time period.

404ErrorRatePct

The percentage of requests for which the HTTP status code was a 404 error for the specified distribution during the time period.

502ErrorRatePct

The percentage of requests for which the HTTP status code was a 502 error for the specified distribution during the time period.

503ErrorRatePct

The percentage of requests for which the HTTP status code was a 503 error for the specified distribution during the time period.

504ErrorRatePct

The percentage of requests for which the HTTP status code was a 504 error for the specified distribution during the time period.

OriginLatency

The total time spent, in milliseconds, from when CloudFront received a request to when it started providing a response to the network (not the viewer), for requests that were served from the origin, not the CloudFront cache. This is also known as *first byte latency*, or *time-to-first-byte*.

CacheHitRate

The percentage of all cacheable requests for which CloudFront served the content from its cache. HTTP POST and PUT requests, and errors, are not considered cacheable requests.

Getting metrics using the CloudWatch API

You can use the Amazon CloudWatch API or CLI to get the CloudFront metrics in programs or applications that you build. You can use the raw data to build your own custom dashboards, your own alarming tools, and so on. To get the CloudFront metrics from the CloudWatch API, you must use the US East (N. Virginia) Region (us-east-1). You also need to know certain values and types for each metric.

Topics

- [Values for all CloudFront metrics \(p. 435\)](#)
- [Values for individual CloudFront metrics \(p. 436\)](#)

Values for all CloudFront metrics

The following values apply to all CloudFront metrics:

Namespace

The value for Namespace is always `AWS/CloudFront`.

Dimensions

Each CloudFront metric has the following two dimensions:

DistributionId

The value for DistributionId is the ID of the CloudFront distribution for which you want to get metrics.

Region

The value for Region is always `Global`, because CloudFront is a global service.

Note

To get the CloudFront metrics from the CloudWatch API, you must use the US East (N. Virginia) Region (us-east-1).

Values for individual CloudFront metrics

Use information from the following list to get details about specific CloudFront metrics from the CloudWatch API. Some of these metrics are available only when you have enabled additional metrics for the distribution.

Note

Only one statistic, *Average* or *Sum*, is applicable for each metric. The following list specifies which statistic is applicable to that metric.

4xx error rate

The percentage of all viewer requests for which the response's HTTP status code is 4xx.

- Metric name: `4xxErrorRate`
- Valid statistic: *Average*
- Unit: *Percent*

401 error rate

The percentage of all viewer requests for which the response's HTTP status code is 401. To get this metric, you must first [enable additional metrics \(p. 430\)](#).

- Metric name: `401ErrorRate`
- Valid statistic: *Average*
- Unit: *Percent*

403 error rate

The percentage of all viewer requests for which the response's HTTP status code is 403. To get this metric, you must first [enable additional metrics \(p. 430\)](#).

- Metric name: `403ErrorRate`
- Valid statistic: *Average*
- Unit: *Percent*

404 error rate

The percentage of all viewer requests for which the response's HTTP status code is 404. To get this metric, you must first [enable additional metrics \(p. 430\)](#).

- Metric name: `404ErrorRate`
- Valid statistic: *Average*
- Unit: *Percent*

5xx error rate

The percentage of all viewer requests for which the response's HTTP status code is 5xx.

- Metric name: `5xxErrorRate`
- Valid statistic: *Average*
- Unit: *Percent*

502 error rate

The percentage of all viewer requests for which the response's HTTP status code is 502. To get this metric, you must first [enable additional metrics \(p. 430\)](#).

- Metric name: `502ErrorRate`
- Valid statistic: *Average*
- Unit: *Percent*

503 error rate

The percentage of all viewer requests for which the response's HTTP status code is 503. To get this metric, you must first [enable additional metrics \(p. 430\)](#).

- Metric name: `503ErrorRate`
- Valid statistic: `Average`
- Unit: `Percent`

504 error rate

The percentage of all viewer requests for which the response's HTTP status code is 504. To get this metric, you must first [enable additional metrics \(p. 430\)](#).

- Metric name: `504ErrorRate`
- Valid statistic: `Average`
- Unit: `Percent`

Bytes downloaded

The total number of bytes downloaded by viewers for GET, HEAD, and OPTIONS requests.

- Metric name: `BytesDownloaded`
- Valid statistic: `Sum`
- Unit: `None`

Bytes uploaded

The total number of bytes that viewers uploaded to your origin with CloudFront, using POST and PUT requests.

- Metric name: `BytesUploaded`
- Valid statistic: `Sum`
- Unit: `None`

Cache hit rate

The percentage of all cacheable requests for which CloudFront served the content from its cache. HTTP POST and PUT requests, and errors, are not considered cacheable requests. To get this metric, you must first [enable additional metrics \(p. 430\)](#).

- Metric name: `CacheHitRate`
- Valid statistic: `Average`
- Unit: `Percent`

Origin latency

The total time spent, in milliseconds, from when CloudFront receives a request to when it starts providing a response to the network (not the viewer), for requests that are served from the origin, not the CloudFront cache. This is also known as *first byte latency*, or *time-to-first-byte*. To get this metric, you must first [enable additional metrics \(p. 430\)](#).

- Metric name: `OriginLatency`
- Valid statistic: `Percentile`
- Unit: `Milliseconds`

Note

To get a `Percentile` statistic from the CloudWatch API, use the `ExtendedStatistics` parameter, not `Statistics`. For more information, see [GetMetricStatistics](#) in the *Amazon CloudWatch API Reference*, or the reference documentation for the [AWS SDKs](#).

Requests

The total number of viewer requests received by CloudFront, for all HTTP methods and for both HTTP and HTTPS requests.

- Metric name: `Requests`

- Valid statistic: Sum
- Unit: None

Total error rate

The percentage of all viewer requests for which the response's HTTP status code is 4xx or 5xx.

- Metric name: `TotalErrorRate`
- Valid statistic: Average
- Unit: Percent

CloudFront logging

Amazon CloudFront provides different kinds of logging. You can log the requests that come to your CloudFront distributions, or you can log the CloudFront service activity in your AWS account.

Logging requests

CloudFront provides the following ways to log the requests that come to your distributions.

Standard logs (access logs)

CloudFront standard logs provide detailed records about every request that's made to a distribution. These logs are useful for many scenarios, including security and access audits.

CloudFront standard logs are delivered to the Amazon S3 bucket of your choice. CloudFront doesn't charge for standard logs, though you incur Amazon S3 charges for storing and accessing the log files.

For more information, see [Using Standard Logs \(Access Logs\)](#) (p. 439).

Real-time logs

CloudFront real-time logs provide information about requests made to a distribution, in real time (log records are delivered within seconds of receiving the requests). You can choose the *sampling rate* for your real-time logs—that is, the percentage of requests for which you want to receive real-time log records. You can also choose the specific fields that you want to receive in the log records.

CloudFront real-time logs are delivered to the data stream of your choice in Amazon Kinesis Data Streams. CloudFront charges for real-time logs, in addition to the charges you incur for using Kinesis Data Streams.

For more information, see [Real-time logs](#) (p. 457).

Logging service activity

You can use AWS CloudTrail to log the CloudFront service activity in your AWS account. CloudTrail provides a record of actions taken by a user, role, or AWS service in CloudFront. Using the information collected by CloudTrail, you can determine the request that was made to CloudFront, the IP address from which the request was made, who made the request, when it was made, and additional details.

For more information, see [Capturing API Requests with CloudTrail](#) (p. 466).

Topics

- [Configuring and Using Standard Logs \(Access Logs\)](#) (p. 439)
- [Real-time logs](#) (p. 457)
- [Using AWS CloudTrail to Capture Requests Sent to the CloudFront API](#) (p. 466)

Configuring and Using Standard Logs (Access Logs)

You can configure CloudFront to create log files that contain detailed information about every user request that CloudFront receives. These are called *standard logs*, also known as *access logs*. These standard logs are available for both web and RTMP distributions. If you enable standard logs, you can also specify the Amazon S3 bucket that you want CloudFront to save files in.

You can enable standard logs when you create or update a distribution. For more information, see [Values That You Specify When You Create or Update a Distribution](#) (p. 38).

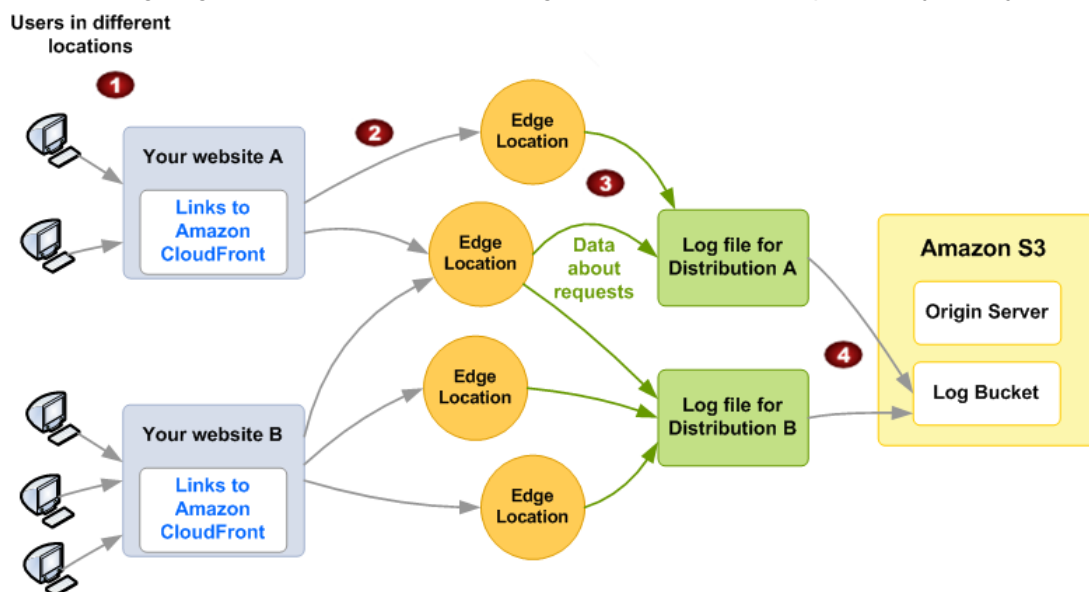
CloudFront also offers real-time logs, which give you information about requests made to a distribution in real time (logs are delivered within seconds of receiving the requests). You can use real-time logs to monitor, analyze, and take action based on content delivery performance. For more information, see [Real-time logs](#) (p. 457).

Topics

- [How Standard Logging Works](#) (p. 439)
- [Choosing an Amazon S3 Bucket for Your Standard Logs](#) (p. 440)
- [Permissions Required to Configure Standard Logging and to Access Your Log Files](#) (p. 440)
- [Required CMK Key Policy for Use with SSE-KMS Buckets](#) (p. 441)
- [File Name Format](#) (p. 442)
- [Timing of Standard Log File Delivery](#) (p. 442)
- [How Requests Are Logged When the Request URL or Headers Exceed the Maximum Size](#) (p. 442)
- [Analyzing Standard Logs](#) (p. 443)
- [Editing Your Standard Logging Settings](#) (p. 443)
- [Deleting Standard Log Files from an Amazon S3 Bucket](#) (p. 443)
- [Standard Log File Format](#) (p. 444)
- [Charges for Standard Logs](#) (p. 456)

How Standard Logging Works

The following diagram shows how CloudFront logs information about requests for your objects.



The following explains how CloudFront logs information about requests for your objects, as illustrated in the previous diagram.

1. In this diagram, you have two websites, A and B, and two corresponding CloudFront distributions. Users request your objects using URLs that are associated with your distributions.
2. CloudFront routes each request to the appropriate edge location.
3. CloudFront writes data about each request to a log file specific to that distribution. In this example, information about requests related to Distribution A goes into a log file just for Distribution A, and information about requests related to Distribution B goes into a log file just for Distribution B.
4. CloudFront periodically saves the log file for a distribution in the Amazon S3 bucket that you specified when you enabled logging. CloudFront then starts saving information about subsequent requests in a new log file for the distribution.

If no users access your content during a given hour, you don't receive any log files for that hour.

Each entry in a log file gives details about a single request. For more information about log file format, see [Standard Log File Format \(p. 444\)](#).

Important

We recommend that you use the logs to understand the nature of the requests for your content, not as a complete accounting of all requests. CloudFront delivers access logs on a best-effort basis. The log entry for a particular request might be delivered long after the request was actually processed and, in rare cases, a log entry might not be delivered at all. When a log entry is omitted from access logs, the number of entries in the access logs won't match the usage that appears in the AWS usage and billing reports.

Choosing an Amazon S3 Bucket for Your Standard Logs

When you enable logging for a distribution, you specify the Amazon S3 bucket that you want CloudFront to store log files in. If you're using Amazon S3 as your origin, we recommend that you don't use the same bucket for your log files; using a separate bucket simplifies maintenance.

Note

Don't choose an Amazon S3 bucket in any of the following Regions, because CloudFront doesn't deliver access logs to buckets in these Regions:

- Africa (Cape Town) af-south-1
- Asia Pacific (Hong Kong) ap-east-1
- Europe (Milan) eu-south-1
- Middle East (Bahrain) me-south-1

The [Amazon S3 console](#) shows the bucket's Region.

You can store the log files for multiple distributions in the same bucket. When you enable logging, you can specify an optional prefix for the file names, so you can keep track of which log files are associated with which distributions.

Permissions Required to Configure Standard Logging and to Access Your Log Files

Your AWS account must have the following permissions for the bucket that you specify for log files:

- The S3 access control list (ACL) for the bucket must grant you `FULL_CONTROL`. If you're the bucket owner, your account has this permission by default. If you're not, the bucket owner must update the ACL for the bucket.

- `s3:GetBucketAcl`
- `s3:PutBucketAcl`

Note the following:

ACL for the bucket

When you create or update a distribution and enable logging, CloudFront uses these permissions to update the ACL for the bucket to give the `awslogsdelivery` account `FULL_CONTROL` permission. The `awslogsdelivery` account writes log files to the bucket. If your account doesn't have the required permissions to update the ACL, creating or updating the distribution will fail.

In some circumstances, if you programmatically submit a request to create a bucket but a bucket with the specified name already exists, S3 resets permissions on the bucket to the default value. If you configured CloudFront to save access logs in an S3 bucket and you stop getting logs in that bucket, check permissions on the bucket to ensure that CloudFront has the necessary permissions.

Restoring the ACL for the bucket

If you remove permissions for the `awslogsdelivery` account, CloudFront won't be able to save logs to the S3 bucket. To enable CloudFront to start saving logs for your distribution again, restore the ACL permission by doing one of the following:

- Disable logging for your distribution in CloudFront, and then enable it again. For more information, see [Values That You Specify When You Create or Update a Distribution \(p. 38\)](#).
- Add the ACL permission for `awslogsdelivery` manually by navigating to the S3 bucket in the Amazon S3 console and adding permission. To add the ACL for `awslogsdelivery`, you must provide the canonical ID for the account, which is the following:

```
c4c1ede66af53448b93c283ce9448c4ba468c9432aa01d700d3878632f77d2d0
```

For more information about adding ACLs to S3 buckets, see [How Do I Set ACL Bucket Permissions?](#) in the *Amazon Simple Storage Service Console User Guide*.

ACL for each log file

In addition to the ACL on the bucket, there's an ACL on each log file. The bucket owner has `FULL_CONTROL` permission on each log file, the distribution owner (if different from the bucket owner) has no permission, and the `awslogsdelivery` account has read and write permissions.

Disabling logging

If you disable logging, CloudFront doesn't delete the ACLs for either the bucket or the log files. If you want, you can do that yourself.

Required CMK Key Policy for Use with SSE-KMS Buckets

If you enabled server-side encryption for your Amazon S3 bucket using AWS KMS-managed keys (SSE-KMS) with a customer-managed Customer Master Key (CMK), you must add the following to the key policy for your CMK to enable writing log files to the bucket. You cannot use the default CMK because CloudFront won't be able to upload the log files to the bucket.

```
{
  "Sid": "Allow CloudFront Flow Logs to use the key",
  "Effect": "Allow",
  "Principal": {
    "Service": "delivery.logs.amazonaws.com"
  },
  "Action": "kms:GenerateDataKey*",
```

```
"Resource": "*"
}
```

File Name Format

The name of each log file that CloudFront saves in your Amazon S3 bucket uses the following file name format:

`<optional prefix>/<distribution ID>.YYYY-MM-DD-HH.unique-ID.gz`

The date and time are in Coordinated Universal Time (UTC).

For example, if you use `example-prefix` as the prefix, and your distribution ID is `EMLARXS9EXAMPLE`, your file names look similar to this:

`example-prefix/EMLARXS9EXAMPLE.2019-11-14-20.RT4KCN4SGK9.gz`

When you enable logging for a distribution, you can specify an optional prefix for the file names, so you can keep track of which log files are associated with which distributions. If you include a value for the log file prefix and your prefix doesn't end with a forward slash (/), CloudFront appends one automatically. If your prefix does end with a forward slash, CloudFront doesn't add another one.

The `.gz` at the end of the file name indicates that CloudFront has compressed the log file using gzip.

Timing of Standard Log File Delivery

CloudFront delivers standard logs for a distribution up to several times an hour. In general, a log file contains information about the requests that CloudFront received during a given time period. CloudFront usually delivers the log file for that time period to your Amazon S3 bucket within an hour of the events that appear in the log. Note, however, that some or all log file entries for a time period can sometimes be delayed by up to 24 hours. When log entries are delayed, CloudFront saves them in a log file for which the file name includes the date and time of the period in which the requests occurred, not the date and time when the file was delivered.

When creating a log file, CloudFront consolidates information for your distribution from all of the edge locations that received requests for your objects during the time period that the log file covers.

CloudFront can save more than one file for a time period depending on how many requests CloudFront receives for the objects associated with a distribution.

CloudFront begins to reliably deliver access logs about four hours after you enable logging. You might get a few access logs before that time.

Note

If no users request your objects during the time period, you don't receive any log files for that period.

CloudFront also offers real-time logs, which give you information about requests made to a distribution in real time (logs are delivered within seconds of receiving the requests). You can use real-time logs to monitor, analyze, and take action based on content delivery performance. For more information, see [Real-time logs \(p. 457\)](#).

How Requests Are Logged When the Request URL or Headers Exceed the Maximum Size

If the total size of all request headers, including cookies, exceeds 20 KB, or if the URL exceeds 8192 bytes, CloudFront can't parse the request completely and can't log the request. Because the request isn't logged, you won't see in the log files the HTTP error status code returned.

If the request body exceeds the maximum size, the request is logged, including the HTTP error status code.

Analyzing Standard Logs

Because you can receive multiple access logs per hour, we recommend that you combine all the log files you receive for a given time period into one file. You can then analyze the data for that period more accurately and completely.

One way to analyze your access logs is to use [Amazon Athena](#). Athena is an interactive query service that can help you analyze data for AWS services, including CloudFront. To learn more, see [Querying Amazon CloudFront Logs](#) in the *Amazon Athena User Guide*.

In addition, the following AWS blog posts discuss some ways to analyze access logs.

- [Amazon CloudFront Request Logging](#) (for content delivered via HTTP)
- [Amazon CloudFront Now Supports Streaming Access Logs](#) (for content delivered via RTMP)
- [Enhanced CloudFront Logs, Now With Query Strings](#)

Important

We recommend that you use the logs to understand the nature of the requests for your content, not as a complete accounting of all requests. CloudFront delivers access logs on a best-effort basis. The log entry for a particular request might be delivered long after the request was actually processed and, in rare cases, a log entry might not be delivered at all. When a log entry is omitted from access logs, the number of entries in the access logs won't match the usage that appears in the AWS usage and billing reports.

Editing Your Standard Logging Settings

You can enable or disable logging, change the Amazon S3 bucket where your logs are stored, and change the prefix for log files by using the [CloudFront console](#) or the CloudFront API. Your changes to logging settings take effect within 12 hours.

For more information, see the following topics:

- To update a web or RTMP distribution using the CloudFront console, see [Updating a Distribution \(p. 63\)](#).
- To update a web distribution using the CloudFront API, see [UpdateDistribution](#) in the *Amazon CloudFront API Reference*.
- To update an RTMP distribution using the CloudFront API, see [UpdateStreamingDistribution](#) in the *Amazon CloudFront API Reference*.

To use the CloudFront API to change access log settings for web distributions, you must use API version 2009-04-02 or later. To use the CloudFront API to change access log settings for RTMP distributions, you must use API version 2010-05-01 or later.

Deleting Standard Log Files from an Amazon S3 Bucket

CloudFront does not automatically delete log files from your Amazon S3 bucket. For information about deleting log files from an Amazon S3 bucket, see the following topics:

- Using the Amazon S3 console: [Deleting Objects](#) in the *Amazon Simple Storage Service Console User Guide*.
- Using the REST API: [DeleteObject](#) in the *Amazon Simple Storage Service API Reference*.

Standard Log File Format

Topics

- [Web Distribution Standard Log File Format \(p. 445\)](#)
- [RTMP Distribution Log File Format \(p. 455\)](#)

Each entry in a log file gives details about a single viewer request. The log files for web and for RTMP distributions are not identical, but they share the following characteristics:

- Use the [W3C extended log file format](#).
- Contain tab-separated values.
- Contain records that are not necessarily in chronological order.
- Contain two header lines: one with the file format version, and another that lists the W3C fields included in each record.
- Contain URL-encoded equivalents for spaces and certain other characters in field values.

URL-encoded equivalents are used for the following characters:

- ASCII character codes 0 through 32, inclusive
- ASCII character codes 127 and higher
- All characters in the following table

The URL encoding standard is defined in [RFC 1738](#).

URL-Encoded Value	Character
%3C	<
%3E	>
%22	"
%23	#
%25	%
%7B	{
%7D	}
%7C	
%5C	\
%5E	^
%7E	~
%5B	[
%5D]
%60	`
%27	'
%20	space

Web Distribution Standard Log File Format

The log file for a web distribution includes the following fields in the listed order.

Field Number	Field Name	Description
1	date	The date on which the event occurred in the format YYYY-MM-DD. For example, 2019-06-30. The date and time are in Coordinated Universal Time (UTC). For WebSocket connections, this is the date when the connection closed.
2	time	The time when the CloudFront server finished responding to the request (in UTC), for example, 01:42:39. For WebSocket connections, this is the time when the connection is closed.
3	x-edge-location	The edge location that served the request. Each edge location is identified by a three-letter code and an arbitrarily assigned number, for example, DFW3. The three-letter code typically corresponds with the International Air Transport Association airport code for an airport near the edge location. (These abbreviations might change in the future.) For a list of edge locations, see the Amazon CloudFront Infrastructure page.
4	sc-bytes	The total number of bytes that CloudFront served to the viewer in response to the request, including headers, for example, 1045619. For WebSocket connections, this is the total number of bytes sent from the server to the client through the connection.
5	c-ip	The IP address of the viewer that made the request, for example, 192.0.2.183 or 2001:0db8:85a3:0000:0000:8a2e:0370:7334. If the viewer used an HTTP proxy or a load balancer to send the request, the value of c-ip is the IP address of the proxy or load balancer. See also X-Forwarded-For in field 20.
6	cs-method	The HTTP request method: DELETE, GET, HEAD, OPTIONS, PATCH, POST, or PUT.
7	cs(Host)	The domain name of the CloudFront distribution, for example, d111111abcdef8.cloudfront.net.
8	cs-uri-stem	The portion of the URI that identifies the path and object, for example, /images/cat.jpg. Question marks (?) in URLs and query strings are not included in the log.
9	sc-status	One of the following values: <ul style="list-style-type: none">An HTTP status code (for example, 200). For a list of HTTP status codes, see RFC 2616 section 10. For more information, see How CloudFront Processes and Caches HTTP 4xx and 5xx Status Codes from Your Origin (p. 286).000, which indicates that the viewer closed the connection (for example, closed the browser tab) before CloudFront could respond to a request. If the viewer closes the connection after CloudFront starts to send the response, the log contains the applicable HTTP status code.

Field Number	Field Name	Description
10	cs(Referer)	The name of the domain that originated the request. Common referrers include search engines, other websites that link directly to your objects, and your own website.
11	cs(User-Agent)	The value of the User-Agent header in the request. The User-Agent header identifies the source of the request, such as the type of device and browser that submitted the request and, if the request came from a search engine, which search engine. For more information, see User-Agent Header (p. 278) .
12	cs-uri-query	<p>The query string portion of the URI, if any. When a URI doesn't contain a query string, this field's value is a hyphen (-).</p> <p>For more information, see Caching Content Based on Query String Parameters (p. 241).</p>
13	cs(Cookie)	<p>The cookie header in the request, including name-value pairs and the associated attributes. If you enable cookie logging, CloudFront logs the cookies in all requests regardless of which cookies you choose to forward to the origin. When a request doesn't include a cookie header, this field's value is a hyphen (-).</p> <p>For more information about cookies, see Caching Content Based on Cookies (p. 244).</p>

Field Number	Field Name	Description
14	x-edge-result-type	<p>How CloudFront classifies the response after the last byte left the edge location. In some cases, the result type can change between the time that CloudFront is ready to send the response and the time that CloudFront has finished sending the response. See also x-edge-response-result-type in field 23.</p> <p>For example, in HTTP streaming, suppose CloudFront finds a segment in the edge cache. In that scenario, the value of this field would ordinarily be <code>Hit</code>. However, if the viewer closes the connection before CloudFront has delivered the entire segment, the final result type, and thus the value of this field, is <code>Error</code>.</p> <p>As another example, WebSocket connections will have a value of <code>Miss</code> for this field because the content is not cacheable and is proxied directly back to the origin server.</p> <p>Possible values include:</p> <ul style="list-style-type: none"> • <code>Hit</code> – CloudFront served the object to the viewer from the edge cache. • <code>RefreshHit</code> – CloudFront found the object in the edge cache but it had expired, so CloudFront contacted the origin to determine whether the cache has the latest version of the object and, if not, to get the latest version. • <code>Miss</code> – The request could not be satisfied by an object in the edge cache, so CloudFront forwarded the request to the origin server and returned the result to the viewer. • <code>LimitExceeded</code> – The request was denied because a CloudFront quota (formerly referred to as a limit) was exceeded. • <code>CapacityExceeded</code> – CloudFront returned an HTTP 503 status code (Service Unavailable) because the CloudFront edge server was temporarily unable to respond to requests. • <code>Error</code> – Typically, this means the request resulted in a client error (<code>sc-status</code> is 4xx) or a server error (<code>sc-status</code> is 5xx). If <code>sc-status</code> is 200, it means the HTTP request was successful but the client disconnected before they downloaded all of the bytes. <p>If the value of this field is <code>Error</code> and the value of x-edge-response-result-type is not <code>Error</code>, it means the client disconnected before finishing the download.</p> <ul style="list-style-type: none"> • <code>Redirect</code> – CloudFront redirects from HTTP to HTTPS. <p>If <code>sc-status</code> is 403 and you configured CloudFront to restrict the geographic distribution of your content, the request might have come from a restricted location. For more information about geo restriction, see Restricting the Geographic Distribution of Your Content (p. 217).</p>

Field Number	Field Name	Description
15	<code>x-edge-request-id</code>	An encrypted string that uniquely identifies a request. In the response header, this is <code>x-amz-cf-id</code> .
16	<code>x-host-header</code>	<p>The value that the viewer included in the <code>Host</code> header for this request. This is the domain name in the request:</p> <ul style="list-style-type: none"> If you're using the CloudFront domain name in your object URLs, such as <code>http://d1111111abcdef8.cloudfront.net/logo.png</code>, this field contains that domain name. If you're using alternate domain names in your object URLs, such as <code>http://example.com/logo.png</code>, this field contains the alternate domain name, such as <i>example.com</i>. To use alternate domain names, you must add them to your distribution. For more information, see Using Custom URLs for Files by Adding Alternate Domain Names (CNAMEs) (p. 71). <p>If you're using alternate domain names, see <code>cs(Host)</code> in field 7 for the domain name that is associated with your distribution.</p>
17	<code>cs-protocol</code>	The protocol that the viewer specified in the request: <code>http</code> , <code>https</code> , <code>ws</code> , or <code>wss</code> .
18	<code>cs-bytes</code>	The number of bytes of data that the viewer included in the request, including headers. For WebSocket connections, this is the total number of bytes sent from the client to the server on the connection.
19	<code>time-taken</code>	The number of seconds (to the thousandth of a second, for example, 0.002) between the time that a CloudFront edge server receives a viewer's request and the time that CloudFront writes the last byte of the response to the edge server's output queue as measured on the server. From the perspective of the viewer, the total time to get the full object will be longer than this value due to network latency and TCP buffering.
20	<code>x-forwarded-for</code>	<p>If the viewer used an HTTP proxy or a load balancer to send the request, the value of <code>c-ip</code> in field 5 is the IP address of the proxy or load balancer. In that case, this field is the IP address of the viewer that originated the request. This field contains IPv4 addresses (such as 192.0.2.44) and IPv6 addresses (such as 2001:0db8:85a3:0000:0000:8a2e:0370:7334), as applicable.</p> <p>If the viewer did not use an HTTP proxy or a load balancer, the value of <code>x-forwarded-for</code> is a hyphen (-).</p>

Field Number	Field Name	Description
21	ssl-protocol	<p>When <code>cs-protocol</code> in field 17 is <code>https</code>, this field contains the SSL/TLS protocol that the client and CloudFront negotiated for transmitting the request and response.</p> <p>Possible values include the following:</p> <ul style="list-style-type: none"> • <code>SSLv3</code> • <code>TLSv1</code> • <code>TLSv1.1</code> • <code>TLSv1.2</code> <p>When <code>cs-protocol</code> in field 17 is <code>http</code>, the value for this field is a hyphen (-).</p>
22	ssl-cipher	<p>When <code>cs-protocol</code> in field 17 is <code>https</code>, this field contains the SSL/TLS cipher that the client and CloudFront negotiated for encrypting the request and response.</p> <p>Possible values include the following:</p> <ul style="list-style-type: none"> • <code>ECDHE-RSA-AES128-GCM-SHA256</code> • <code>ECDHE-RSA-AES128-SHA256</code> • <code>ECDHE-RSA-AES128-SHA</code> • <code>ECDHE-RSA-AES256-GCM-SHA384</code> • <code>ECDHE-RSA-AES256-SHA384</code> • <code>ECDHE-RSA-AES256-SHA</code> • <code>AES128-GCM-SHA256</code> • <code>AES256-GCM-SHA384</code> • <code>AES128-SHA256</code> • <code>AES256-SHA</code> • <code>AES128-SHA</code> • <code>DES-CBC3-SHA</code> • <code>RC4-MD5</code> <p>When <code>cs-protocol</code> in field 17 is <code>http</code>, the value for this field is a hyphen (-).</p>

Field Number	Field Name	Description
23	x-edge-response-result-type	<p>How CloudFront classified the response just before returning the response to the viewer. See also <code>x-edge-result-type</code> in field 14.</p> <p>Possible values include:</p> <ul style="list-style-type: none">• <code>Hit</code> – CloudFront served the object to the viewer from the edge cache.• <code>RefreshHit</code> – CloudFront found the object in the edge cache but it had expired, so CloudFront contacted the origin to verify that the cache has the latest version of the object.• <code>Miss</code> – The request could not be satisfied by an object in the edge cache, so CloudFront forwarded the request to the origin server and returned the result to the viewer.• <code>LimitExceeded</code> – The request was denied because a CloudFront quota (formerly referred to as a limit) was exceeded.• <code>CapacityExceeded</code> – CloudFront returned a 503 error because the edge location didn't have enough capacity at the time of the request to serve the object.• <code>Error</code> – Typically, this means the request resulted in a client error (<code>sc-status</code> is 4xx) or a server error (<code>sc-status</code> is 5xx). <p>If the value of <code>x-edge-result-type</code> is <code>Error</code> and the value of this field is not <code>Error</code>, the client disconnected before finishing the download.</p> <ul style="list-style-type: none">• <code>Redirect</code> – CloudFront redirects from HTTP to HTTPS. <p>If <code>sc-status</code> is 403 and you configured CloudFront to restrict the geographic distribution of your content, the request might have come from a restricted location. For more information about geo restriction, see Restricting the Geographic Distribution of Your Content (p. 217).</p>
24	cs-protocol-version	<p>The HTTP version that the viewer specified in the request. Possible values include:</p> <ul style="list-style-type: none">• <code>HTTP/0.9</code>• <code>HTTP/1.0</code>• <code>HTTP/1.1</code>• <code>HTTP/2.0</code>

Field Number	Field Name	Description
25	<code>file-status</code>	<p>When field-level encryption is configured for a distribution, this field contains a code that indicates whether the request body was successfully processed. If field-level encryption is not configured for the distribution, the value of this field is a hyphen (-).</p> <p>When CloudFront successfully processes the request body, encrypts values in the specified fields, and forwards the request to the origin, the value of this field is <code>Processed</code>. The value of <code>x-edge-result-type</code>, field 14, can still indicate a client-side or server-side error in this case.</p> <p>If the request exceeds a field-level encryption quota, <code>file-status</code> contains one of the following error codes, and CloudFront returns HTTP status code 400 to the viewer. For a list of the current quotas on field-level encryption, see Quotas on Field-Level Encryption (p. 500).</p> <ul style="list-style-type: none"> <code>FieldLengthLimitClientError</code> – A field that is configured to be encrypted exceeded the maximum length allowed. <code>FieldNumberLimitClientError</code> – A request that CloudFront is configured to encrypt contains more than the number of fields allowed. <code>RequestLengthLimitClientError</code> – The length of the request body exceeded the maximum length allowed when field-level encryption is configured. <p>Other possible values for this field include the following:</p> <ul style="list-style-type: none"> <code>ForwardedByContentType</code> – CloudFront forwarded the request to the origin without parsing or encryption because no content type was configured. <code>ForwardedByQueryArgs</code> – CloudFront forwarded the request to the origin without parsing or encryption because the request contains a query argument that wasn't in the configuration for field-level encryption. <code>ForwardedDueToNoProfile</code> – CloudFront forwarded the request to the origin without parsing or encryption because no profile was specified in the configuration for field-level encryption. <code>MalformedContentTypeClientError</code> – CloudFront rejected the request and returned an HTTP 400 status code to the viewer because the value of the <code>Content-Type</code> header was in an invalid format. <code>MalformedInputClientError</code> – CloudFront rejected the request and returned an HTTP 400 status code to the viewer because the request body was in an invalid format. <code>MalformedQueryArgsClientError</code> – CloudFront rejected the request and returned an HTTP 400 status code to the viewer because a query argument was empty or in an invalid format. <code>RejectedByContentType</code> – CloudFront rejected the request and returned an HTTP 400 status code to the viewer because no

Field Number	Field Name	Description
		content type was specified in the configuration for field-level encryption. <ul style="list-style-type: none">• <code>RejectedByQueryArgs</code> – CloudFront rejected the request and returned an HTTP 400 status code to the viewer because no query argument was specified in the configuration for field-level encryption.• <code>ServerError</code> – The server returned an error.
26	<code>file-encrypted-fields</code>	The number of fields that CloudFront encrypted and forwarded to the origin. CloudFront streams the processed request to the origin as it encrypts data, so <code>file-encrypted-fields</code> can have a value even if the value of <code>file-status</code> is an error. If field-level encryption is not configured for the distribution, the value of <code>file-encrypted-fields</code> is a hyphen (-).
27	<code>c-port</code>	The port number of the request from the viewer.
28	<code>time-to-first-byte</code>	The number of seconds between receiving the request and writing the first byte of the response, as measured on the server.

Field Number	Field Name	Description
29	x-edge-detailed-result-type	<p>When x-edge-result-type (field 14) is not <code>Error</code>, this field contains the same value as x-edge-result-type. When x-edge-result-type is <code>Error</code>, this field contains the specific type of error.</p> <p>Possible error type values for this field include the following:</p> <ul style="list-style-type: none"> • <code>AbortedOrigin</code> – CloudFront encountered an issue with the origin. • <code>ClientCommError</code> – The response to the viewer was interrupted due to a communication problem between CloudFront and the viewer. • <code>ClientGeoBlocked</code> – The distribution is configured to refuse requests from the viewer's geographic location. • <code>ClientHungUpRequest</code> – The viewer stopped prematurely while sending the request. • <code>Error</code> – An error occurred for which the error type doesn't fit any of the other categories. This error type can occur when CloudFront serves an error response from the CloudFront cache. • <code>InvalidRequest</code> – CloudFront received an invalid request from the viewer. • <code>InvalidRequestBlocked</code> – Access to the requested resource is blocked. • <code>InvalidRequestCertificate</code> – The distribution doesn't match the SSL/TLS certificate for which the HTTPS connection was established. • <code>InvalidRequestHeader</code> – The request contained an invalid header field or value. • <code>InvalidRequestMethod</code> – The distribution is not configured to handle the HTTP request method that was used. This can happen when the distribution supports only cacheable requests. • <code>OriginConnectError</code> – CloudFront couldn't connect to the origin. • <code>OriginContentRangeLengthError</code> – The <code>Content-Length</code> header in the origin's response doesn't match the length in the <code>Content-Range</code> header. • <code>OriginDnsError</code> – CloudFront couldn't resolve the origin's domain name. • <code>OriginError</code> – The origin returned an incorrect response. • <code>OriginHeaderTooBigError</code> – A header returned by the origin is too big for CloudFront to process. • <code>OriginInvalidResponseError</code> – The origin returned an invalid response. • <code>OriginReadError</code> – CloudFront couldn't read from the origin. • <code>OriginWriteError</code> – CloudFront couldn't write to the origin. • <code>OriginZeroSizeObjectError</code> – A zero size object sent from the origin resulted in an error. • <code>SlowReaderOriginError</code> – The viewer was slow to read the message that caused the origin error.

Field Number	Field Name	Description
30	sc-content-type	The value of the HTTP Content-Type header of the response.
31	sc-content-len	The value of the HTTP Content-Length header of the response.
32	sc-range-start	When the response contains the HTTP Content-Range header, this field contains the range start value.
33	sc-range-end	When the response contains the HTTP Content-Range header, this field contains the range end value.

The following is an example log file for a web distribution:

```
#Version: 1.0
#Fields: date time x-edge-location sc-bytes c-ip cs-method cs(Host) cs-uri-stem sc-status
cs(Referer) cs(User-Agent) cs-uri-query cs(Cookie) x-edge-result-type x-edge-request-id
x-host-header cs-protocol cs-bytes time-taken x-forwarded-for ssl-protocol ssl-cipher x-
edge-response-result-type cs-protocol-version file-status file-encrypted-fields c-port time-
to-first-byte x-edge-detailed-result-type sc-content-type sc-content-len sc-range-start sc-
range-end
2019-12-04 21:02:31 LAX1 392 192.0.2.100 GET d111111abcdef8.cloudfront.net /index.html
200 - Mozilla/5.0%20(Windows%20NT%2010.0;%20Win64;%20x64)%20AppleWebKit/537.36%20(KHTML,
%20like%20Gecko)%20Chrome/78.0.3904.108%20Safari/537.36 - - Hit
SOX4xwn4XV6Q4rgb7XiVGOHms_BG1TAC4KyHmureZmBNrjGdRLiNIQ== d111111abcdef8.cloudfront.net
https 23 0.001 - TLSv1.2 ECDHE-RSA-AES128-GCM-SHA256 Hit HTTP/2.0 - - 11040 0.001 Hit
text/html 78 - -
2019-12-04 21:02:31 LAX1 392 192.0.2.100 GET d111111abcdef8.cloudfront.net /index.html
200 - Mozilla/5.0%20(Windows%20NT%2010.0;%20Win64;%20x64)%20AppleWebKit/537.36%20(KHTML,
%20like%20Gecko)%20Chrome/78.0.3904.108%20Safari/537.36 - - Hit
k5WGMNkEzR5BEM_SaF47gjtX9zBD02m349OY2an0QPEaUum1ZOLrow== d111111abcdef8.cloudfront.net
https 23 0.000 - TLSv1.2 ECDHE-RSA-AES128-GCM-SHA256 Hit HTTP/2.0 - - 11040 0.000 Hit
text/html 78 - -
2019-12-04 21:02:31 LAX1 392 192.0.2.100 GET d111111abcdef8.cloudfront.net /index.html
200 - Mozilla/5.0%20(Windows%20NT%2010.0;%20Win64;%20x64)%20AppleWebKit/537.36%20(KHTML,
%20like%20Gecko)%20Chrome/78.0.3904.108%20Safari/537.36 - - Hit
f37nTMVvnKvV2ZSvEsivup_c2kZ7VXzYdjC-GUQZ5qNs-89BlWazbw== d111111abcdef8.cloudfront.net
https 23 0.001 - TLSv1.2 ECDHE-RSA-AES128-GCM-SHA256 Hit HTTP/2.0 - - 11040 0.001 Hit
text/html 78 - -
2019-12-13 22:36:27 SEA19-C1 900 192.0.2.200 GET d111111abcdef8.cloudfront.net /
favicon.ico 502 http://www.example.com/ Mozilla/5.0%20(Windows
%20NT%2010.0;%20Win64;%20x64)%20AppleWebKit/537.36%20(KHTML,
%20like%20Gecko)%20Chrome/78.0.3904.108%20Safari/537.36 - - Error
1pkpNfBQ39sYmnpjUQjmH2w1wdJnbHYTbag21o_30fcQgPzdL2RSSQ== www.example.com http 675 0.102 -
- - Error HTTP/1.1 - - 25260 0.102 OriginDnsError text/html 507 - -
2019-12-13 22:36:26 SEA19-C1 900 192.0.2.200 GET d111111abcdef8.cloudfront.net / 502
- Mozilla/5.0%20(Windows%20NT%2010.0;%20Win64;%20x64)%20AppleWebKit/537.36%20(KHTML,
%20like%20Gecko)%20Chrome/78.0.3904.108%20Safari/537.36 - - Error
3AqrZGcNf_g0-5KovfA7c9XLcf4YGvMFSeFdiEtRiN_2y8jSis8Zxg== www.example.com http 735 0.107 -
- - Error HTTP/1.1 - - 3802 0.107 OriginDnsError text/html 507 - -
2019-12-13 22:37:02 SEA19-C2 900 192.0.2.200 GET d111111abcdef8.cloudfront.net / 502
- curl/7.55.1 - - Error kBkDzGnceVtWHqSCqBUqtA_cEs2T3tFUBbnBNkB9El_uVRhHgcZfcw==
www.example.com http 387 0.103 - - - Error HTTP/1.1 - - 12644 0.103 OriginDnsError text/
html 507 - -
```

RTMP Distribution Log File Format

Each record in an RTMP access log represents a playback event, for example, connect, play, pause, stop, disconnect, and so on. As a result, CloudFront generates multiple log records each time a viewer watches a video. To relate log records that stem from the same stream ID, use the `x-sid` field.

Note

Some fields have values for all events, and some have values only for play, stop, pause, unpause, and seek events. Usually, when the log file contains a hyphen (-) for a field, it means that field isn't relevant for the corresponding event.

The following table describes the fields that are present in each record in the RTMP distribution log file, regardless of the type of event. The fields appear in the log in the order listed.

Field Number	Field Name	Description
1	<code>date</code>	The date on which the event occurred in the format YYYY-MM-DD, for example, 2019-05-23. The date and time are in Coordinated Universal Time (UTC).
2	<code>time</code>	The time when the server received the request (in UTC), for example, 01:42:39.
3	<code>x-edge-location</code>	The edge location where the playback event occurred. Each edge location is identified by a three-letter code and an arbitrarily assigned number, for example, DFW3. The three-letter code typically corresponds with the International Air Transport Association airport code for an airport near the edge location. (These abbreviations might change in the future.) For a list of edge locations, see the Amazon CloudFront Infrastructure page.
4	<code>c-ip</code>	Client IP, for example, 192.0.2.183.
5	<code>x-event</code>	The event type. This is a Connect, Disconnect, Play, Stop, Pause, Unpause, or Seek event.
6	<code>sc-bytes</code>	The running total number of bytes sent from the server to the client, up to the time of the event.
7	<code>x-cf-status</code>	A code indicating the status of the event. Currently, "OK" is the only value for this field. New functionality in the future could require new status codes.
8	<code>x-cf-client-id</code>	An opaque string identifier that can be used to differentiate clients. This value is unique for each connection.
9	<code>cs-uri-stem</code>	The stem portion of the URI, including the application and the application instance. This is sometimes referred to as the FMS <i>connect string</i> . For example, <code>rtmp://shqshne4jdp4b6.cloudfront.net/cfx/st</code> . Question marks (?) in URLs and query strings are not included in the log.
10	<code>cs-uri-query</code>	The query string portion of the URI that is included on the connect string. Question marks (?) in URLs and query strings are not included in the log.
11	<code>c-referrer</code>	The URI of the referrer.
12	<code>x-page-url</code>	The URL of the page from which the SWF is linked.

Field Number	Field Name	Description
13	c-user-agent	The value of the User-Agent header in the request. The User-Agent header identifies the type of device that submitted the request. For more information, see User-Agent Header (p. 278) .

The following fields usually have values only for Play, Stop, Pause, Unpause, and Seek events. For other events, they contain a single hyphen (-). These fields appear in the log after the fields in the preceding table and in the order listed.

Field Number	Field Name	Description
14	x-sname	The stream name.
15	x-sname-query	The stream query string, if any.
16	x-file-ext	The stream type, for example, FLV.
17	x-sid	The stream ID. This is a unique integer identifier for the connection.

The following is an example of a log file for an RTMP distribution:

```
#Version: 1.0
#Fields: date time x-edge-location c-ip x-event sc-bytes x-cf-status x-cf-client-id cs-uri-
stem cs-uri-query c-referrer x-page-url c-user-agent x-sname x-sname-query x-file-ext x-sid
2010-03-12 23:51:20 SEA4 192.0.2.147 connect 2014 OK
bfd8a98bee0840d9b871b7f6ade9908f rtmp://shqshne4jdp4b6.cloudfront.net/cfx/st key=value
http://player.example.com/player.swf http://www.example.com/support/jw-player-setup-
wizard?example=204 LNX%2010,0,32,18 - - - -
2010-03-12 23:51:21 SEA4 192.0.2.222 play 3914 OK bfd8a98bee0840d9b871b7f6ade9908f
rtmp://shqshne4jdp4b6.cloudfront.net/cfx/st key=value http://player.example.com/
player.swf http://www.example.com/support/jw-player-setup-wizard?example=204 LNX
%2010,0,32,18 myvideo p=2&q=4 flv 1
2010-03-12 23:53:44 SEA4 192.0.2.4 stop 323914 OK bfd8a98bee0840d9b871b7f6ade9908f
rtmp://shqshne4jdp4b6.cloudfront.net/cfx/st key=value http://player.example.com/
player.swf http://www.example.com/support/jw-player-setup-wizard?example=204 LNX
%2010,0,32,18 dir/other/myvideo p=2&q=4 flv 1
2010-03-12 23:53:44 SEA4 192.0.2.103 play 8783724 OK
bfd8a98bee0840d9b871b7f6ade9908f rtmp://shqshne4jdp4b6.cloudfront.net/cfx/st key=value
http://player.example.com/player.swf http://www.example.com/support/jw-player-setup-
wizard?example=204 LNX%2010,0,32,18 dir/favs/myothervideo p=42&q=14 mp4 2
2010-03-12 23:56:21 SEA4 192.0.2.199 stop 429822014 OK
bfd8a98bee0840d9b871b7f6ade9908f rtmp://shqshne4jdp4b6.cloudfront.net/cfx/st key=value
http://player.example.com/player.swf http://www.example.com/support/jw-player-setup-
wizard?example=204 LNX%2010,0,32,18 dir/favs/myothervideo p=42&q=14 mp4 2
2010-03-12 23:59:44 SEA4 192.0.2.14 disconnect 429824092 OK
bfd8a98bee0840d9b871b7f6ade9908f rtmp://shqshne4jdp4b6.cloudfront.net/cfx/st key=value
http://player.example.com/player.swf http://www.example.com/support/jw-player-setup-
wizard?example=204 LNX%2010,0,32,18 - - - -
```

Charges for Standard Logs

Standard logging is an optional feature of CloudFront. There is no extra charge for enabling standard logging. However, you accrue the usual Amazon S3 charges for storing and accessing the files on Amazon S3 (you can delete them at any time).

For more information about Amazon S3 pricing, see [Amazon S3 Pricing](#).

For more information about CloudFront pricing, see [CloudFront Pricing](#).

Real-time logs

With CloudFront real-time logs, you can get information about requests made to a distribution in real time (logs are delivered within seconds of receiving the requests). You can use real-time logs to monitor, analyze, and take action based on content delivery performance.

CloudFront real-time logs are configurable. You can choose:

- The *sampling rate* for your real-time logs—that is, the percentage of requests for which you want to receive real-time log records.
- The specific fields that you want to receive in the log records.
- The specific cache behaviors (path patterns) that you want to receive real-time logs for.

CloudFront real-time logs are delivered to the data stream of your choice in Amazon Kinesis Data Streams. You can build your own [Kinesis data stream consumer](#), or use Amazon Kinesis Data Firehose to send the log data to Amazon Simple Storage Service (Amazon S3), Amazon Redshift, Amazon Elasticsearch Service (Amazon ES), or a third-party log processing service.

CloudFront charges for real-time logs, in addition to the charges you incur for using Kinesis Data Streams. For more information about pricing, see [Amazon CloudFront Pricing](#) and [Amazon Kinesis Data Streams pricing](#).

Understanding real-time log configurations

To use CloudFront real-time logs, you start by creating a real-time log configuration. The real-time log configuration contains information about which log fields you want to receive, the *sampling rate* for log records, and the Kinesis data stream where you want to deliver the logs.

Specifically, a real-time log configuration contains the following settings:

- [Name](#) (p. 457)
- [Sampling rate](#) (p. 457)
- [Fields](#) (p. 457)
- [Endpoint \(Kinesis data stream\)](#) (p. 461)
- [IAM role](#) (p. 462)

Name

A name to identify the real-time log configuration.

Sampling rate

The sampling rate is a whole number between 1 and 100 (inclusive) that determines the percentage of viewer requests that are sent to Kinesis Data Streams as real-time log records. To include every viewer request in your real-time logs, specify 100 for the sampling rate. You might choose a lower sampling rate to reduce costs while still receiving a representative sample of request data in your real-time logs.

Fields

A list of the fields that are included in each real-time log record. You can choose to receive all of the available fields, or only the fields that you need for monitoring and analyzing performance.

The following list describes each of the available fields. The fields are listed in the order in which they appear in the log records that are delivered to Kinesis Data Streams.

- **timestamp** – The date and time at which the edge server finished responding to the request.
- **c-ip** – The IP address of the viewer that made the request, for example, 192.0.2.183 or 2001:0db8:85a3:0000:0000:8a2e:0370:7334. If the viewer used an HTTP proxy or a load balancer to send the request, the value of this field is the IP address of the proxy or load balancer.
- **time-to-first-byte** – The number of seconds between receiving the request and writing the first byte of the response, as measured on the edge server.
- **sc-status** – Contains one of the following values:
 - The HTTP status code of the edge server’s response (for example, 200).
 - 000, which indicates that the viewer closed the connection before the edge server could respond to the request. If the viewer closes the connection after the edge server starts to send the response, this field contains the applicable HTTP status code.
- **sc-bytes** – The total number of bytes that the edge server sent to the viewer in response to the request, including headers.
- **cs-method** – The HTTP request method received from the viewer.
- **cs-protocol** – The protocol of the viewer request (http, https, ws, or wss).
- **cs-host** – The domain name of the CloudFront distribution (for example, d111111abcdef8.cloudfront.net).
- **cs-uri-stem** – The portion of the request URL that identifies the path and object (for example, /images/cat.jpg). Question marks (?) in URLs and query strings are not included in the log.
- **cs-bytes** – The total number of bytes of data that the viewer included in the request, including headers.
- **x-edge-location** – The edge location that served the request. Each edge location is identified by a three-letter code and an arbitrarily assigned number (for example, DFW3). The three-letter code typically corresponds with the International Air Transport Association (IATA) airport code for an airport near the edge location’s geographic location. (These abbreviations might change in the future.)
- **x-edge-request-id** – An opaque string that uniquely identifies a request. This string is also sent in the x-amz-cf-id response header.
- **x-host-header** – The value that the viewer included in the Host header of the request. If you’re using the CloudFront domain name in your object URLs (such as d111111abcdef8.cloudfront.net), this field contains that domain name. If you’re using alternate domain names (CNAMEs) in your object URLs (such as www.example.com), this field contains the alternate domain name.
- **time-taken** – The number of seconds (to the thousandth of a second, for example, 0.082) from when the edge server receives the viewer’s request to when the edge server writes the last byte of the response to the output queue, as measured on the edge server. From the perspective of the viewer, the total time to get the full response will be longer than this value because of network latency and TCP buffering.
- **cs-protocol-version** – The HTTP version that the viewer specified in the request. Possible values include HTTP/0.9, HTTP/1.0, HTTP/1.1, and HTTP/2.0.
- **c-ip-version** – The IP version of the request (IPv4 or IPv6).
- **cs-user-agent** – The value of the User-Agent header in the request. The User-Agent header identifies the source of the request, such as the type of device and browser that submitted the request or, if the request came from a search engine, which search engine.
- **cs-referrer** – The value of the Referer header in the request. This is the name of the domain that originated the request. Common referrers include search engines, other websites that link directly to your objects, and your own website.
- **cs-cookie** – The Cookie header in the request, including name—value pairs and the associated attributes.

Note

This field is truncated to 800 bytes.

- **cs-uri-query** – The query string portion of the URL, if any.
- **x-edge-response-result-type** – How the edge server classified the response just before returning the response to the viewer. Possible values include:
 - **Hit** – The edge server served the object to the viewer from the cache.
 - **RefreshHit** – The edge server found the object in the edge cache but the object had expired, so the edge server contacted the origin to verify that the cache had the latest version of the object.
 - **Miss** – The request could not be satisfied by an object in the edge cache, so the edge server forwarded the request to the origin server and returned the result to the viewer.
 - **LimitExceeded** – The request was denied because a CloudFront quota (formerly referred to as a limit) was exceeded.
 - **CapacityExceeded** – The edge server returned a 503 error because it didn't have enough capacity at the time of the request to serve the object.
 - **Error** – Typically, this means the request resulted in a client error or a server error.
 - **Redirect** – The edge server redirected the viewer from HTTP to HTTPS according to the distribution settings.
- **x-forwarded-for** – If the viewer used an HTTP proxy or a load balancer to send the request, the value of the `c-ip` field is the IP address of the proxy or load balancer. In that case, this field is the IP address of the viewer that originated the request.
- **ssl-protocol** – When the request used HTTPS, this field contains the SSL/TLS protocol that the viewer and edge server negotiated for transmitting the request and response. For a list of possible values, see the supported SSL/TLS protocols in [Supported protocols and ciphers between viewers and CloudFront \(p. 128\)](#).
- **ssl-cipher** – When the request used HTTPS, this field contains the SSL/TLS cipher that the viewer and edge server negotiated for encrypting the request and response. For a list of possible values, see the supported SSL/TLS protocols in [Supported protocols and ciphers between viewers and CloudFront \(p. 128\)](#).
- **x-edge-result-type** – How the edge server classified the response after the last byte left the edge location. In some cases, the result type can change between the time that the edge server is ready to send the response and the time that it finishes sending the response. (See also the `x-edge-response-result-type` field).

For example, in HTTP streaming, suppose the edge server finds a segment of the stream in the edge cache. In that scenario, the value of this field would ordinarily be `Hit`. However, if the viewer closes the connection before the edge server has delivered the entire segment, the final result type (and the value of this field) is `Error`.

Possible values include:

- **Hit** – The edge server served the object to the viewer from the cache.
- **RefreshHit** – The edge server found the object in the edge cache but the object had expired, so the edge server contacted the origin to verify that the cache had the latest version of the object.
- **Miss** – The request could not be satisfied by an object in the edge cache, so the edge server forwarded the request to the origin server and returned the result to the viewer.
- **LimitExceeded** – The request was denied because a CloudFront quota (formerly referred to as a limit) was exceeded.
- **CapacityExceeded** – The edge server returned a 503 error because it didn't have enough capacity at the time of the request to serve the object.
- **Error** – Typically, this means the request resulted in a client error or a server error.
- **Redirect** – The edge server redirected the viewer from HTTP to HTTPS according to the distribution settings.
- **file-encrypted-fields** – The number of [field-level encryption \(p. 219\)](#) fields that the edge server encrypted and forwarded to the origin. Edge servers stream the processed request to the origin as they encrypt data, so this field can have a value even if the value of `file-status` is an error.

- **file-status** – When [field-level encryption](#) is configured for a distribution, this field contains a code that indicates whether the request body was successfully processed. When the edge server successfully processes the request body, encrypts values in the specified fields, and forwards the request to the origin, the value of this field is `Processed`. The value of `x-edge-result-type` can still indicate a client-side or server-side error in this case.

If the request exceeds a field-level encryption quota (formerly referred to as a limit), this field contains one of the following error codes, and the edge server returns HTTP status code 400 to the viewer.

- **FieldLengthLimitClientError** – A field that is configured to be encrypted exceeded the maximum length allowed.
- **FieldNumberLimitClientError** – A request that the distribution is configured to encrypt contains more than the number of fields allowed.
- **RequestLengthLimitClientError** – The length of the request body exceeded the maximum length allowed when field-level encryption is configured.

Other possible values for this field include the following:

- **ForwardedByContentType** – The edge server forwarded the request to the origin without parsing or encryption because no content type was configured.
- **ForwardedByQueryArgs** – The edge server forwarded the request to the origin without parsing or encryption because the request contains a query argument that wasn't in the configuration for field-level encryption.
- **ForwardedDueToNoProfile** – The edge server forwarded the request to the origin without parsing or encryption because no profile was specified in the configuration for field-level encryption.
- **MalformedContentTypeClientError** – The edge server rejected the request and returned an HTTP 400 status code to the viewer because the value of the `Content-Type` header was in an invalid format.
- **MalformedInputClientError** – The edge server rejected the request and returned an HTTP 400 status code to the viewer because the request body was in an invalid format.
- **MalformedQueryArgsClientError** – The edge server rejected the request and returned an HTTP 400 status code to the viewer because a query argument was empty or in an invalid format.
- **RejectedByContentType** – The edge server rejected the request and returned an HTTP 400 status code to the viewer because no content type was specified in the configuration for field-level encryption.
- **RejectedByQueryArgs** – The edge server rejected the request and returned an HTTP 400 status code to the viewer because no query argument was specified in the configuration for field-level encryption.
- **ServerError** – The origin server returned an error.
- **sc-content-type** – The value of the HTTP `Content-Type` header of the response.
- **sc-content-len** – The value of the HTTP `Content-Length` header of the response.
- **sc-range-start** – When the response contains the HTTP `Content-Range` header, this field contains the range start value.
- **sc-range-end** – When the response contains the HTTP `Content-Range` header, this field contains the range end value.
- **c-port** – The port number of the request from the viewer.
- **x-edge-detailed-result-type** – When the `x-edge-result-type` field is not `Error`, this field contains the same value as `x-edge-result-type`. When `x-edge-result-type` is `Error`, this field contains the specific type of error. Possible values include:
 - **AbortedOrigin** – The edge server encountered an issue with the origin.
 - **ClientCommError** – The response to the viewer was interrupted due to a communication problem between the edge server and the viewer.
 - **ClientGeoBlocked** – The distribution is configured to refuse requests from the viewer's geographic location.

- **ClientHungUpRequest** – The viewer stopped prematurely while sending the request.
- **Error** – An error occurred for which the error type doesn't fit any of the other categories. This error type can occur when the edge server serves an error response from the cache.
- **InvalidRequest** – The edge server received an invalid request from the viewer.
- **InvalidRequestBlocked** – Access to the requested resource is blocked.
- **InvalidRequestCertificate** – The distribution doesn't match the SSL/TLS certificate for which the HTTPS connection was established.
- **InvalidRequestHeader** – The request contained an invalid header.
- **InvalidRequestMethod** – The distribution is not configured to handle the HTTP request method that was used. This can happen when the distribution supports only cacheable requests.
- **OriginConnectError** – The edge server couldn't connect to the origin.
- **OriginContentRangeLengthError** – The Content-Length header in the origin's response doesn't match the length in the Content-Range header.
- **OriginDnsError** – The edge server couldn't resolve the origin's domain name.
- **OriginError** – The origin returned an incorrect response.
- **OriginHeaderTooBigError** – A header returned by the origin is too big for the edge server to process.
- **OriginInvalidResponseError** – The origin returned an invalid response.
- **OriginReadError** – The edge server couldn't read from the origin.
- **OriginWriteError** – The edge server couldn't write to the origin.
- **OriginZeroSizeObjectError** – A zero size object sent from the origin resulted in an error.
- **SlowReaderOriginError** – The viewer was slow to read the message that caused the origin error.
- **c-country** – A country code that represents the viewer's geographic location, as determined by the viewer's IP address.
- **cs-accept-encoding** – The value of the Accept-Encoding header in the viewer request.
- **cs-accept** – The value of the Accept header in the viewer request.
- **cache-behavior-path-pattern** – The path pattern that identifies the cache behavior that matched the viewer request.
- **cs-headers** – The HTTP headers (names and values) in the viewer request.

Note
This field is truncated to 800 bytes.
- **cs-header-names** – The names of the HTTP headers (not values) in the viewer request.

Note
This field is truncated to 800 bytes.
- **cs-headers-count** – The number of HTTP headers in the viewer request.

Endpoint (Kinesis data stream)

The endpoint contains information about the Kinesis data stream where you want to send real-time logs. You provide the Amazon Resource Name (ARN) of the data stream.

For more information about creating a Kinesis data stream, see the following topics in the *Amazon Kinesis Data Streams Developer Guide*.

- [Managing Streams Using the Console](#)
- [Perform Basic Kinesis Data Stream Operations Using the AWS CLI](#)
- [Creating a Stream](#) (uses the AWS SDK for Java)

When you create a data stream, you need to specify the number of shards. Use the following information to help you estimate the number of shards you need.

To estimate the number of shards for your Kinesis data stream

1. Calculate (or estimate) the number of requests per second that your CloudFront distribution receives.

You can use the [CloudFront usage reports](#) (in the CloudFront console) and the [CloudFront metrics \(p. 429\)](#) (in the CloudFront and Amazon CloudWatch consoles) to help you calculate your requests per second.

2. Determine the typical size of a single real-time log record.

In general, a single log record is around 500 bytes. A large record that includes all available fields is around 1 KB.

If you're not sure what your log record size is, you can enable real-time logs with a low sampling rate (for example, 1%), and then calculate the average record size using monitoring data in Kinesis Data Streams (total number of records divided by total incoming bytes).

3. Multiply the number of requests per second (from step 1) by the size of a typical real-time log record (from step 2) to determine the amount of data per second that your real-time log configuration is likely to send to the Kinesis data stream.
4. Using the data per second, calculate the number of shards that you need. A single shard can handle no more than 1 MB per second and 1,000 requests (log records) per second. When calculating the number of shards that you need, we recommend adding 10-25% as a buffer.

For example, assume your distribution receives 50,000 requests per second, and that your real-time log records size is typically 500 bytes. This means that your real-time log configuration could generate 25,000,000 bytes (50,000 multiplied by 500), or 23.84 MB, per second. In this scenario you would need at least 24 shards. To add a buffer of about 20%, you would specify 29 shards.

IAM role

The AWS Identity and Access Management (IAM) role that gives CloudFront permission to deliver real-time logs to your Kinesis data stream.

When you create a real-time log configuration with the CloudFront console, you can choose **Create new service role** to let the console create the IAM role for you.

When you create a real-time log configuration with AWS CloudFormation or the CloudFront API (AWS CLI or SDK), you must create the IAM role yourself and provide the role ARN. To create the IAM yourself, use the following policies.

IAM role trust policy

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "cloudfront.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

IAM role permissions policy for an unencrypted data stream

To use the following policy, replace *Kinesis data stream ARN* with the ARN of your Kinesis data stream.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "kinesis:DescribeStreamSummary",
        "kinesis:DescribeStream",
        "kinesis:PutRecord",
        "kinesis:PutRecords"
      ],
      "Resource": [
        "Kinesis data stream ARN"
      ]
    }
  ]
}
```

IAM role permissions policy for an encrypted data stream

To use the following policy, replace *Kinesis data stream ARN* with the ARN of your Kinesis data stream and *AWS KMS key* with the ARN of your customer master key in AWS Key Management Service (AWS KMS).

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "kinesis:DescribeStreamSummary",
        "kinesis:DescribeStream",
        "kinesis:PutRecord",
        "kinesis:PutRecords"
      ],
      "Resource": [
        "Kinesis data stream ARN"
      ]
    },
    {
      "Effect": "Allow",
      "Action": [
        "kms:GenerateDataKey"
      ],
      "Resource": [
        "AWS KMS key"
      ]
    }
  ]
}
```

Creating and using real-time log configurations

You can use a real-time log configurations to get information about requests made to a distribution in real time (logs are delivered within seconds of receiving the requests). You can create a real-time log configuration in the CloudFront console, with the AWS Command Line Interface (AWS CLI), or with the CloudFront API.

To use a real-time log configuration, you attach it to one or more cache behaviors in a CloudFront distribution.

Create a real-time log configuration (console)

To create a real-time log configuration (console)

1. Sign in to the AWS Management Console and open the **Logs** page in the CloudFront console at <https://console.aws.amazon.com/cloudfront/v2/home?#/logs>.
2. Choose **Real-time log configurations**.
3. Choose **Create configuration**.
4. Choose the desired setting for the real-time log configuration. Note the following:
 - By default, all **Fields** are chosen. To remove a field, do one of the following:
 - Use the **Choose fields** drop-down menu to remove the selection from the fields that you don't want to include in the real-time log configuration.
 - Use the expand button (⊕) to view all fields, then use the remove button (✕) to remove the fields that you don't want to include in the real-time log configuration.
 - For **IAM role**, you can choose **Create new service role** to let the console create the IAM role for you. You must have permission to create IAM roles.
 - You can use the setting in the **Distribution** section to choose a CloudFront distribution and cache behavior to attach to the real-time log configuration.

For more information, see [Understanding real-time log configurations \(p. 457\)](#).

5. When finished, choose **Create configuration**.

If successful, the console shows the details of the real-time log configuration that you just created.

Create a real-time log configuration (AWS CLI)

To create a real-time log configuration with the AWS Command Line Interface (AWS CLI), use the **aws cloudfront create-realtime-log-config** command. You can use an input file to provide the command's input parameters, rather than specifying each individual parameter as command line input.

To create a real-time log configuration (CLI with input file)

1. Use the following command to create a file named `rtl-config.yaml` that contains all of the input parameters for the **create-realtime-log-config** command.

```
aws cloudfront create-realtime-log-config --generate-cli-skeleton yaml-input > rtl-config.yaml
```

Note

The `yaml-input` option is available only in [version 2 of the AWS CLI](#). With version 1 of the AWS CLI, you can generate an input file in JSON format. For more information, see [Generating AWS CLI skeleton and input parameters from a JSON or YAML input file](#) in the *AWS Command Line Interface User Guide*.

2. Open the file named `rtl-config.yaml` that you just created. Edit the file to specify the real-time log configuration settings that you want, then save the file. Note the following:
 - For `StreamType`, the only valid value is `Kinesis`.

For more information about the real-time log configuration settings, see [Understanding real-time log configurations \(p. 457\)](#).

3. Use the following command to create the real-time log configuration using input parameters from the `rtl-config.yaml` file.

```
aws cloudfront create-realtime-log-config --cli-input-yaml file://rtl-config.yaml
```

If successful, the command's output shows the details of the real-time log configuration that you just created.

To attach a real-time log configuration to an existing distribution (CLI with input file)

1. Use the following command to save the distribution configuration for the CloudFront distribution that you want to update. Replace `distribution_ID` with the distribution's ID.

```
aws cloudfront get-distribution-config --id distribution_ID --output yaml > dist-config.yaml
```

Note

The `--output yaml` option is available only in [version 2 of the AWS CLI](#). With version 1 of the AWS CLI, you can generate the output in JSON format. For more information, see [Controlling command output from the AWS CLI](#) in the *AWS Command Line Interface User Guide*.

2. Open the file named `dist-config.yaml` that you just created. Edit the file, making the following changes to each cache behavior that you are updating to use a real-time log configuration.
 - In the cache behavior, add a field named `RealtimeLogConfigArn`. For the field's value, use the ARN of the real-time log configuration that you want to attach to this cache behavior.
 - Rename the `ETag` field to `IfMatch`, but don't change the field's value.

Save the file when finished.

3. Use the following command to update the distribution to use the real-time log configuration. Replace `distribution_ID` with the distribution's ID.

```
aws cloudfront update-distribution --id distribution_ID --cli-input-yaml file://dist-config.yaml
```

If successful, the command's output shows the details of the distribution that you just updated.

Create a real-time log configuration (API)

To create a real-time log configuration with the CloudFront API, use [CreateRealtimeLogConfig](#). For more information about the parameters that you specify in this API call, see [Understanding real-time log configurations \(p. 457\)](#) and the API reference documentation for your AWS SDK or other API client.

After you create a real-time log configuration, you can attach it to a cache behavior, using one of the following API calls:

- To attach it to a cache behavior in an existing distribution, use [UpdateDistribution](#).
- To attach it to a cache behavior in a new distribution, use [CreateDistribution](#).

For both of these API calls, provide the ARN of the real-time log configuration in the `RealtimeLogConfigArn` field, inside a cache behavior. For more information about the other fields that you specify in these API calls, see [Values That You Specify When You Create or Update a Distribution](#) (p. 38) and the API reference documentation for your AWS SDK or other API client.

Creating a Kinesis Data Streams consumer

To read and analyze your real-time logs, you build or use a Kinesis Data Streams *consumer*. When you build a consumer for CloudFront real-time logs, it's important to know that the fields in every real-time log record are always delivered in the same order, as listed in the [Fields](#) (p. 457) section. Make sure that you build your consumer to accommodate this fixed order.

For example, consider a real-time log configuration that includes only these three fields: `time-to-first-byte`, `sc-status`, and `c-country`. In this scenario, the last field, `c-country`, is always field number 3 in every log record. However, if you later add fields to the real-time log configuration, the placement of each field in a record can change.

For example, if you add the fields `sc-bytes` and `time-taken` to the real-time log configuration, these fields are inserted into each log record according to the order shown in the [Fields](#) (p. 457) section. The resulting order of all five fields is `time-to-first-byte`, `sc-status`, `sc-bytes`, `time-taken`, and `c-country`. The `c-country` field was originally field number 3, but is now field number 5. Make sure that your consumer application can handle fields that change position in a log record, in case you add fields to your real-time log configuration.

Troubleshooting real-time logs

After you create a real-time log configuration, you might find that no records (or not all records) are delivered to Kinesis Data Streams. In this case, you should first verify that your CloudFront distribution is receiving viewer requests. If it is, you can check the following setting to continue troubleshooting.

IAM role permissions

To deliver real-time log records to your Kinesis data stream, CloudFront uses the IAM role in the real-time log configuration. Make sure that the role trust policy and the role permissions policy match the policies shown in [IAM role](#) (p. 462).

Kinesis Data Streams throttling

If CloudFront writes real-time log records to your Kinesis data stream faster than the stream can handle, Kinesis Data Streams might throttle the requests from CloudFront. In this case, you can increase the number of shards in your Kinesis data stream. Each shard can support writes up to 1,000 records per second, up to a maximum data write of 1 MB per second.

Using AWS CloudTrail to Capture Requests Sent to the CloudFront API

CloudFront is integrated with CloudTrail, an AWS service that captures information about every request that is sent to the CloudFront API by your AWS account, including your IAM users. CloudTrail periodically saves log files of these requests to an Amazon S3 bucket that you specify. CloudTrail captures information about all requests, whether they were made using the CloudFront console, the CloudFront API, the AWS SDKs, the CloudFront CLI, or another service, for example, AWS CloudFormation.

You can use information in the CloudTrail log files to determine which requests were made to CloudFront, the source IP address from which each request was made, who made the request, when it was made, and so on. To learn more about CloudTrail, including how to configure and enable it, see the [AWS CloudTrail User Guide](#).

Note

CloudFront is a global service. To view CloudFront requests in CloudTrail logs, you must update an existing trail to include global services. For more information, see [Updating a Trail](#) and [About Global Service Events](#) in the *AWS CloudTrail User Guide*.

Topics

- [CloudFront Information in CloudTrail](#) (p. 467)
- [Understanding CloudFront Log File Entries](#) (p. 467)

CloudFront Information in CloudTrail

CloudTrail is enabled on your AWS account when you create the account. When activity occurs in CloudFront, that activity is recorded in a CloudTrail event along with other AWS service events in **Event history**. You can view, search, and download recent events in your AWS account. Because CloudFront is a global service, events for the service are logged in US East (N. Virginia). For more information, see [Viewing Events with CloudTrail Event History](#).

For an ongoing record of events in your AWS account, including events for CloudFront, create a trail. Your trail must include global service events. A trail enables CloudTrail to deliver log files to an Amazon S3 bucket. By default, when you create a trail in the console, the trail applies to all regions and includes global service events. The trail logs events from all regions in the AWS partition and delivers the log files to the Amazon S3 bucket that you specify. Additionally, you can configure other AWS services to further analyze and act upon the event data collected in CloudTrail logs. For more information, see:

- [Overview for Creating a Trail](#)
- [CloudTrail Supported Services and Integrations](#)
- [Configuring Amazon SNS Notifications for CloudTrail](#)
- [Receiving CloudTrail Log Files from Multiple Regions](#) and [Receiving CloudTrail Log Files from Multiple Accounts](#)

All CloudFront API actions are logged by CloudTrail and are documented in the [Amazon CloudFront API Reference](#). For example, calls to the `CreateDistribution`, `GetDistribution` and `ListInvalidations` APIs generate entries in the CloudTrail log files.

Every event or log entry contains information about who generated the request. The identity information helps you determine the following:

- Whether the request was made with root or IAM user credentials.
- Whether the request was made with temporary security credentials for a role or federated user.
- Whether the request was made by another AWS service.

For more information, see the [CloudTrail userIdentity Element](#).

Understanding CloudFront Log File Entries

Each JSON-formatted CloudTrail log file can contain one or more log entries. A log entry represents a single request from any source and includes information about the requested action, including any parameters, the date and time of the action, and so on. The log entries are not guaranteed to be in any particular order; they are not an ordered stack trace of API calls.

The `eventName` element identifies the action that occurred and the API version that was used to perform that action. For example, the following `eventName` value indicates that a web distribution was updated, and the 2014-01-31 API version was used to perform the action:

UpdateDistribution2014_01_31

The following example shows a CloudTrail log entry that demonstrates five actions:

- Updating a web distribution configuration. The value of `eventName` is `UpdateDistribution`.
- Listing web distributions that are associated with the current account. The value of `eventName` is `ListDistributions`.
- Getting the configuration for a specific web distribution. The value of `eventName` is `GetDistribution`.
- Creating an invalidation batch request. The value of `eventName` is `CreateInvalidation`.
- Listing origin access identities that are associated with the current account. The value of `eventName` is `ListCloudFrontOriginAccessIdentities`.

```
{
  "Records": [{
    "eventVersion": "1.01",
    "userIdentity": {
      "type": "IAMUser",
      "principalId": "A1B2C3D4E5F6G7EXAMPLE",
      "arn": "arn:aws:iam::111122223333:user/smithj",
      "accountId": "111122223333",
      "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
      "userName": "smithj"
    },
    "eventTime": "2014-05-06T18:00:32Z",
    "eventName": "UpdateDistribution2014_01_31",
    "sourceIPAddress": "192.0.2.17",
    "userAgent": "aws-sdk-ruby/1.39.0 ruby/1.9.3 x86_64-linux",
    "requestParameters": {
      "id": "EDFDVBD6EXAMPLE",
      "ifMatch": "E9LHASXEXAMPLE",
      "distributionConfig": {
        "restrictions": {
          "geoRestriction": {
            "quantity": 0,
            "restrictionType": "none"
          }
        },
        "customErrorResponses": {
          "quantity": 0
        },
        "defaultRootObject": "index.html",
        "aliases": {
          "quantity": 1,
          "items": ["example.com"]
        },
        "logging": {
          "bucket": "",
          "enabled": false,
          "prefix": "",
          "includeCookies": false
        },
        "viewerCertificate": {
          "iAMCertificateId": "A1B2C3D4E5F6G7EXAMPLE",
          "sSLSupportMethod": "sni-only"
        },
        "callerReference": "2014-05-06 64832",
        "defaultCacheBehavior": {
          "targetOriginId": "Images",
          "allowedMethods": {
            "items": ["GET",
              "HEAD"],
```

```

        "quantity": 2
    },
    "forwardedValues": {
        "cookies": {
            "forward": "none"
        },
        "queryString": false
    },
    "minTTL": 300,
    "trustedSigners": {
        "enabled": false,
        "quantity": 0
    },
    "viewerProtocolPolicy": "redirect-to-https",
    "smoothStreaming": false
},
"origins": {
    "items": [{
        "customOriginConfig": {
            "hTTPSPort": 443,
            "originProtocolPolicy": "http-only",
            "hTTPPort": 80
        },
        "domainName": "myawsbucket.s3-website-us-east-2.amazonaws.com",
        "id": "Web page origin"
    },
    {
        "customOriginConfig": {
            "hTTPSPort": 443,
            "originProtocolPolicy": "http-only",
            "hTTPPort": 80
        },
        "domainName": "myotherawsbucket.s3-website-us-west-2.amazonaws.com",
        "id": "Images"
    }
    ],
    "quantity": 2
},
"enabled": true,
"cacheBehaviors": {
    "allowedMethods": {
        "items": ["GET",
            "HEAD"],
        "quantity": 2
    },
    "trustedSigners": {
        "enabled": false,
        "quantity": 0
    },
    "targetOriginId": "Web page origin",
    "smoothStreaming": false,
    "viewerProtocolPolicy": "redirect-to-https",
    "minTTL": 300,
    "forwardedValues": {
        "cookies": {
            "forward": "none"
        },
        "queryString": false
    },
    "pathPattern": "*.html"
    },
    "quantity": 1
},
"priceClass": "PriceClass_All",
"comment": "Added an origin and a cache behavior"
}
},

```

```

    "responseElements": {
      "eTag": "E2QWRUHEXAMPLE",
      "distribution": {
        "domainName": "d11111abcdef8.cloudfront.net",
        "status": "InProgress",
        "distributionConfig": {
          distributionConfig response omitted
        },
        "id": "EDFDVBD6EXAMPLE",
        "lastModifiedTime": "May 6, 2014 6:00:32 PM",
        "activeTrustedSigners": {
          "quantity": 0,
          "enabled": false
        },
        "inProgressInvalidationBatches": 0
      }
    },
    "requestID": "4e6b66f9-d548-11e3-a8a9-73e33example",
    "eventID": "5ab02562-0fc5-43d0-b7b6-90293example"
  },
  {
    "eventVersion": "1.01",
    "userIdentity": {
      "type": "IAMUser",
      "principalId": "A1B2C3D4E5F6G7EXAMPLE",
      "arn": "arn:aws:iam::111122223333:user/smithj",
      "accountId": "111122223333",
      "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
      "userName": "smithj"
    },
    "eventTime": "2014-05-06T18:01:35Z",
    "eventName": "ListDistributions2014_01_31",
    "sourceIPAddress": "192.0.2.17",
    "userAgent": "aws-sdk-ruby/1.39.0 ruby/1.9.3 x86_64-linux",
    "requestParameters": null,
    "responseElements": null,
    "requestID": "52de9f97-d548-11e3-8fb9-4dad0example",
    "eventID": "eb91f423-6dd3-4bb0-a148-3cdfbexample"
  },
  {
    "eventVersion": "1.01",
    "userIdentity": {
      "type": "IAMUser",
      "principalId": "A1B2C3D4E5F6G7EXAMPLE",
      "arn": "arn:aws:iam::111122223333:user/smithj",
      "accountId": "111122223333",
      "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
      "userName": "smithj"
    },
    "eventTime": "2014-05-06T18:01:59Z",
    "eventName": "GetDistribution2014_01_31",
    "sourceIPAddress": "192.0.2.17",
    "userAgent": "aws-sdk-ruby/1.39.0 ruby/1.9.3 x86_64-linux",
    "requestParameters": {
      "id": "EDFDVBD6EXAMPLE"
    },
    "responseElements": null,
    "requestID": "497b3622-d548-11e3-8fb9-4dad0example",
    "eventID": "c32289c7-005a-46f7-9801-cba41example"
  },
  {
    "eventVersion": "1.01",
    "userIdentity": {
      "type": "IAMUser",
      "principalId": "A1B2C3D4E5F6G7EXAMPLE",
      "arn": "arn:aws:iam::111122223333:user/smithj",

```

```

        "accountId": "111122223333",
        "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
        "userName": "smithj"
    },
    "eventTime": "2014-05-06T18:02:27Z",
    "eventName": "CreateInvalidation2014_01_31",
    "sourceIPAddress": "192.0.2.17",
    "userAgent": "aws-sdk-ruby/1.39.0 ruby/1.9.3 x86_64-linux",
    "requestParameters": {
        "invalidationBatch": {
            "callerReference": "2014-05-06 64947",
            "paths": {
                "quantity": 3,
                "items": ["/images/new.jpg",
                    "/images/logo.jpg",
                    "/images/banner.jpg"]
            }
        }
    },
    "distributionId": "EDFDVBD6EXAMPLE"
},
"responseElements": {
    "invalidation": {
        "createTime": "May 6, 2014 6:02:27 PM",
        "invalidationBatch": {
            "callerReference": "2014-05-06 64947",
            "paths": {
                "quantity": 3,
                "items": ["/images/banner.jpg",
                    "/images/logo.jpg",
                    "/images/new.jpg"]
            }
        }
    },
    "status": "InProgress",
    "id": "ISRZ85EXAMPLE"
},
"location": "https://cloudfront.amazonaws.com/2014-01-31/distribution/EDFDVBD6EXAMPLE/invalidation/ISRZ85EXAMPLE"
},
"requestID": "4e200613-d548-11e3-a8a9-73e33example",
"eventID": "191ebb93-66b7-4517-a741-92b0eexample"
},
{
    "eventVersion": "1.01",
    "userIdentity": {
        "type": "IAMUser",
        "principalId": "A1B2C3D4E5F6G7EXAMPLE",
        "arn": "arn:aws:iam::111122223333:user/smithj",
        "accountId": "111122223333",
        "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
        "userName": "smithj"
    },
    "eventTime": "2014-05-06T18:03:08Z",
    "eventName": "ListCloudFrontOriginAccessIdentities2014_01_31",
    "sourceIPAddress": "192.0.2.17",
    "userAgent": "aws-sdk-ruby/1.39.0 ruby/1.9.3 x86_64-linux",
    "requestParameters": null,
    "responseElements": null,
    "requestID": "42ca4299-d548-11e3-8fb9-4dad0example",
    "eventID": "7aeb434f-eb55-4e2a-82d8-417d5example"
}
}]
}

```


Tracking Configuration Changes with AWS Config

You can use AWS Config to record configuration changes for CloudFront distribution settings changes. For example, you can capture changes to distribution states, price classes, origins, geo restriction settings, and Lambda@Edge configurations.

Note

AWS Config does not record key-value tags for CloudFront distributions.

Set up AWS Config with CloudFront

When you set up AWS Config, you can choose to record all supported AWS resources, or you can specify only certain resources to record configuration changes for, such as just recording changes for CloudFront. To see the specific resources supported for CloudFront, see the list of [Supported AWS Resource Types](#) in the *AWS Config Developer Guide*.

To track configuration changes to your CloudFront distribution, you must log in to the AWS Console in the US East (N. Virginia) public region.

Note

There might be a delay in recording resources with AWS Config. AWS Config records resources only after it discovers the resources.

Set up AWS Config with CloudFront by using the AWS Management Console

1. Sign in to the AWS Management Console and open the AWS Config console at <https://console.aws.amazon.com/config/>.
2. Choose **Get Started Now**.
3. On the **Settings** page, for **Resource types to record**, specify the AWS resource types that you want AWS Config to record. If you want to record only CloudFront changes, choose **Specific types**, and then, under **CloudFront**, choose the distribution or streaming distribution that you want to track changes for.

To add or change which distributions to track, choose **Settings** on the left, after completing your initial setup.

4. Specify additional required options for AWS Config: set up a notification, specify a location for the configuration information, and add rules for evaluating resource types.

For more information, see [Setting up AWS Config with the Console](#) in the *AWS Config Developer Guide*.

To set up AWS Config with CloudFront by using the AWS CLI or by using an API, see one of the following:

- **Use the AWS CLI:** [Setting up AWS Config with the AWS CLI](#) in the *AWS Config Developer Guide*
- **Use an API:** The [StartConfigurationRecorder](#) action and other information in the *AWS Config API Reference*

View CloudFront Configuration History

After AWS Config starts recording configuration changes to your distributions, you can get the configuration history of any distribution that you have configured for CloudFront.

You can view configuration histories in any of the following ways:

- **Use the AWS Config console.** For each recorded resource, you can view a timeline page, which provides a history of configuration details. To view this page, choose the gray icon in the **Config**

Timeline column of the **Dedicated Hosts** page. For more information, see [Viewing Configuration Details in the AWS Config Console](#) in the *AWS Config Developer Guide*.

- **Run AWS CLI commands.** To get a list of all your distributions, use the [list-discovered-resources](#) command. To get the configuration details of a distribution for a specific time interval, use the [get-resource-config-history](#) command. For more information, see [View Configuration Details Using the CLI](#) in the *AWS Config Developer Guide*.
- **Use the AWS Config API in your applications.** To get a list of all your distributions use the [ListDiscoveredResources](#) action. To get the configuration details of a distribution for a specific time interval, use the [GetResourceConfigHistory](#) action. For more information, see the [AWS Config API Reference](#).

For example, to get a list of all of your distributions from AWS Config, you could run a CLI command such as the following:

```
aws configservice list-discovered-resources --resource-type  
AWS::CloudFront::Distribution
```

Or, to get a list of all of your RTMP streaming distributions from AWS Config, run a CLI command such as the following:

```
aws configservice list-discovered-resources --resource-type  
AWS::CloudFront::StreamingDistribution
```

Security in Amazon CloudFront

Cloud security at AWS is the highest priority. As an AWS customer, you benefit from a data center and network architecture that is built to meet the requirements of the most security-sensitive organizations.

Security is a shared responsibility between AWS and you. The [shared responsibility model](#) describes this as security *of* the cloud and security *in* the cloud:

- **Security of the cloud** – AWS is responsible for protecting the infrastructure that runs AWS services in the AWS Cloud. AWS also provides you with services that you can use securely. Third-party auditors regularly test and verify the effectiveness of our security as part of the [AWS compliance programs](#). To learn about the compliance programs that apply to Amazon CloudFront, see [AWS Services in Scope by Compliance Program](#).
- **Security in the cloud** – Your responsibility is determined by the AWS service that you use. You are also responsible for other factors including the sensitivity of your data, your organization's requirements, and applicable laws and regulations.

This documentation helps you understand how to apply the shared responsibility model when using CloudFront. The following topics show you how to configure CloudFront to meet your security and compliance objectives. You also learn how to use other AWS services that help you to monitor and secure your CloudFront resources.

Topics

- [Data Protection in Amazon CloudFront \(p. 474\)](#)
- [Identity and Access Management in CloudFront \(p. 476\)](#)
- [Logging and Monitoring in Amazon CloudFront \(p. 493\)](#)
- [Compliance Validation for Amazon CloudFront \(p. 493\)](#)
- [Resilience in Amazon CloudFront \(p. 495\)](#)
- [Infrastructure Security in Amazon CloudFront \(p. 495\)](#)

Data Protection in Amazon CloudFront

Amazon CloudFront conforms to the AWS [shared responsibility model](#), which includes regulations and guidelines for data protection. AWS is responsible for protecting the global infrastructure that runs all the AWS services. AWS maintains control over data hosted on this infrastructure, including the security configuration controls for handling customer content and personal data. AWS customers and APN partners, acting either as data controllers or data processors, are responsible for any personal data that they put in the AWS Cloud.

For data protection purposes, we recommend that you protect AWS account credentials and set up individual user accounts with AWS Identity and Access Management (IAM), so that each user is given only the permissions necessary to fulfill their job duties. We also recommend that you secure your data in the following ways:

- Use multi-factor authentication (MFA) with each account.
- Use SSL/TLS to communicate with AWS resources.
- Set up API and user activity logging with AWS CloudTrail.
- Use AWS encryption solutions, along with all default security controls within AWS services.

- Use advanced managed security services such as Amazon Macie, which assists in discovering and securing personal data that is stored in Amazon S3.

We strongly recommend that you never put sensitive identifying information, such as your customers' account numbers, into free-form fields such as a **Name** field. This includes when you work with CloudFront or other AWS services using the console, API, AWS CLI, or AWS SDKs. Any data that you enter into CloudFront or other services might get picked up for inclusion in diagnostic logs. When you provide a URL to an external server, don't include credentials information in the URL to validate your request to that server.

For more information about data protection, see the [AWS Shared Responsibility Model and GDPR](#) blog post on the *AWS Security Blog*.

Amazon CloudFront provides several options that you can use to help secure the content that it delivers:

- Configure HTTPS connections.
- Configure field-level encryption to encrypt data during transit.
- Restrict access to content so that only specific people, or people in a specific area, can view it.

The following topics explain the options in more detail.

Topics

- [Encryption in Transit](#) (p. 475)
- [Encryption at Rest](#) (p. 475)
- [Restrict Access to Content](#) (p. 475)

Encryption in Transit

To encrypt your data during transit, you configure Amazon CloudFront to require that viewers use HTTPS to request your files, so that connections are encrypted when CloudFront communicates with viewers. You also can configure CloudFront to use HTTPS to get files from your origin, so that connections are encrypted when CloudFront communicates with your origin.

For more information, see [Using HTTPS with CloudFront](#) (p. 121).

Field-level encryption adds an additional layer of security along with HTTPS that lets you protect specific data throughout system processing so that only certain applications can see it. By configuring field-level encryption in CloudFront, you can securely upload user-submitted sensitive information to your web servers. The sensitive information provided by your clients is encrypted at the edge closer to the user. It remains encrypted throughout your entire application stack, ensuring that only applications that need the data—and have the credentials to decrypt it—are able to do so.

For more information, see [Using Field-Level Encryption to Help Protect Sensitive Data](#) (p. 219).

Encryption at Rest

CloudFront uses SSDs which are encrypted for edge location points of presence (POPs), and encrypted EBS volumes for Regional Edge Caches (RECs).

Restrict Access to Content

Many companies that distribute content over the internet want to restrict access to documents, business data, media streams, or content that is intended for a subset of users. To securely serve this content by using Amazon CloudFront, you can do one or more of the following:

Use signed URLs or cookies

You can restrict access to content that is intended for selected users—for example, users who have paid a fee—by serving this private content through CloudFront using signed URLs or signed cookies. For more information, see [Serving Private Content with Signed URLs and Signed Cookies \(p. 145\)](#).

Restrict access to content in Amazon S3 buckets

If you restrict access to your content by using, for example, CloudFront signed URLs or signed cookies, you also won't want people to view files by using the direct URL for the file. Instead, you want them to access the files only by using the CloudFront URL, so that your protections work.

If you use an Amazon S3 bucket as the origin for a CloudFront distribution, you can set up an origin access identity (OAI) to manage direct access to your content. An origin access identity is a special CloudFront user identity that you can associate with your distribution so that you can secure all or just some of your Amazon S3 content. For more information about how to configure this, see [Restricting Access to Amazon S3 Content by Using an Origin Access Identity \(p. 209\)](#).

Use AWS WAF web ACLs

You can use AWS WAF, a web application firewall service, to create a web access control list (web ACL) to restrict access to your content. Based on conditions that you specify, such as the IP addresses that requests originate from or the values of query strings, CloudFront responds to requests either with the requested content or with an HTTP 403 status code (Forbidden). For more information, see [Using AWS WAF to Control Access to Your Content \(p. 216\)](#).

Use geo restriction

You can use *geo restriction*, also known as *geoblocking*, to prevent users in specific geographic locations from accessing content that you serve through a CloudFront distribution. There are several options to choose from when you configure geo restrictions. For more information, see [Restricting the Geographic Distribution of Your Content \(p. 217\)](#).

Identity and Access Management in CloudFront

To perform any operation on CloudFront resources, such as creating a distribution or invalidating an object, [AWS Identity and Access Management \(IAM\)](#) requires you to authenticate that you're an approved AWS user. If you're using the CloudFront console, you authenticate your identity by providing your AWS user name and a password. If you're accessing CloudFront programmatically, your application authenticates your identity for you by using access keys or by signing requests.

After you authenticate your identity, IAM controls your access to AWS by verifying that you have permissions to perform operations and access resources. If you are an account administrator, you can use IAM to control the access of other users to the resources that are associated with your account.

This chapter explains how to use [AWS Identity and Access Management \(IAM\)](#) and CloudFront to help secure your resources.

Topics

- [Authentication \(p. 476\)](#)
- [Access Control \(p. 478\)](#)

Authentication

You can access AWS as any of the following types of identities:

- **AWS account root user** – When you first create an AWS account, you begin with a single sign-in identity that has complete access to all AWS services and resources in the account. This identity is called the AWS account *root user* and is accessed by signing in with the email address and password that you used to create the account. We strongly recommend that you do not use the root user for your everyday tasks, even the administrative ones. Instead, adhere to the [best practice of using the root user only to create your first IAM user](#). Then securely lock away the root user credentials and use them to perform only a few account and service management tasks.
- **IAM user** – An [IAM user](#) is an identity within your AWS account that has specific custom permissions (for example, permissions to create a web distribution in CloudFront). You can use an IAM user name and password to sign in to secure AWS webpages like the [AWS Management Console](#), [AWS Discussion Forums](#), or the [AWS Support Center](#).

In addition to a user name and password, you can also generate [access keys](#) for each user. You can use these keys when you access AWS services programmatically, either through [one of the several SDKs](#) or by using the [AWS Command Line Interface \(CLI\)](#). The SDK and CLI tools use the access keys to cryptographically sign your request. If you don't use AWS tools, you must sign the request yourself. CloudFront supports *Signature Version 4*, a protocol for authenticating inbound API requests. For more information about authenticating requests, see [Signature Version 4 Signing Process](#) in the *AWS General Reference*.

- **IAM role** – An [IAM role](#) is an IAM identity that you can create in your account that has specific permissions. An IAM role is similar to an IAM user in that it is an AWS identity with permissions policies that determine what the identity can and cannot do in AWS. However, instead of being uniquely associated with one person, a role is intended to be assumable by anyone who needs it. Also, a role does not have standard long-term credentials such as a password or access keys associated with it. Instead, when you assume a role, it provides you with temporary security credentials for your role session. IAM roles with temporary credentials are useful in the following situations:
- **Federated user access** – Instead of creating an IAM user, you can use existing identities from AWS Directory Service, your enterprise user directory, or a web identity provider. These are known as *federated users*. AWS assigns a role to a federated user when access is requested through an [identity provider](#). For more information about federated users, see [Federated Users and Roles](#) in the *IAM User Guide*.
- **AWS service access** – A service role is an IAM role that a service assumes to perform actions in your account on your behalf. When you set up some AWS service environments, you must define a role for the service to assume. This service role must include all the permissions that are required for the service to access the AWS resources that it needs. Service roles vary from service to service, but many allow you to choose your permissions as long as you meet the documented requirements for that service. Service roles provide access only within your account and cannot be used to grant access to services in other accounts. You can create, modify, and delete a service role from within IAM. For example, you can create a role that allows Amazon Redshift to access an Amazon S3 bucket on your behalf and then load data from that bucket into an Amazon Redshift cluster. For more information, see [Creating a Role to Delegate Permissions to an AWS Service](#) in the *IAM User Guide*.
- **Applications running on Amazon EC2** – You can use an IAM role to manage temporary credentials for applications that are running on an EC2 instance and making AWS CLI or AWS API requests. This is preferable to storing access keys within the EC2 instance. To assign an AWS role to an EC2 instance and make it available to all of its applications, you create an instance profile that is attached to the instance. An instance profile contains the role and enables programs that are running on the

EC2 instance to get temporary credentials. For more information, see [Using an IAM Role to Grant Permissions to Applications Running on Amazon EC2 Instances](#) in the *IAM User Guide*.

Access Control

To create, update, delete, or list CloudFront resources, you need permissions to perform the operation, and you need permissions to access the corresponding resources. In addition, to perform the operation programmatically, you need valid access keys.

The following sections describe how to manage permissions for CloudFront:

- [Overview of Managing Access Permissions to Your CloudFront Resources](#) (p. 478)
- [Using Identity-Based Policies \(IAM Policies\) for CloudFront](#) (p. 483)
- [CloudFront API Permissions: Actions, Resources, and Conditions Reference](#) (p. 487)

Overview of Managing Access Permissions to Your CloudFront Resources

Every AWS resource is owned by an AWS account, and permissions to create or access a resource are governed by permissions policies.

Note

An *account administrator* (or administrator user) is a user that has administrator privileges. For more information about administrators, see [IAM Best Practices](#) in the *IAM User Guide*.

When you grant permissions, you decide who gets the permissions, the resources they get permissions for, and the actions that they get permission to perform.

Topics

- [ARNs for CloudFront Resources](#) (p. 478)
- [Understanding Resource Ownership](#) (p. 478)
- [Managing Access to Resources](#) (p. 479)
- [Specifying Policy Elements: Resources, Actions, Effects, and Principals](#) (p. 482)
- [Specifying Conditions in a Policy](#) (p. 482)

ARNs for CloudFront Resources

All CloudFront resources—web and RTMP distributions, invalidations, and origin access identities—use the same format for Amazon Resource Names (ARNs):

```
arn:aws:cloudfront::optional-account-id:*
```

CloudFront provides API actions to work with each of these types of resources. For more information, see the [Amazon CloudFront API Reference](#). For a list of actions and the ARN that you specify to grant or deny permission to use each action, see [CloudFront API Permissions: Actions, Resources, and Conditions Reference](#) (p. 487).

Understanding Resource Ownership

An AWS account owns the resources that are created in the account, regardless of who created the resources. Specifically, the resource owner is the AWS account of the principal entity (that is, the root account, an IAM user, or an IAM role) that authenticates the resource creation request.

The following examples illustrate how this works:

- If you use the root account credentials of your AWS account to create a web distribution, your AWS account is the owner of the distribution.
- If you create an IAM user in your AWS account and grant permissions to create a web distribution to that user, the user can create a web distribution. The AWS account that created the user owns the distribution.
- If you create an IAM role in your AWS account with permissions to create a web distribution, anyone who can assume the role can create a web distribution. Your AWS account, to which the role belongs, owns the distribution.

Managing Access to Resources

A *permissions policy* specifies who has access to what. This section explains the options for creating permissions policies for CloudFront. For general information about IAM policy syntax and descriptions, see the [AWS IAM Policy Reference](#) in the *IAM User Guide*.

Policies attached to an IAM identity are referred to as identity-based policies (IAM policies), and policies attached to a resource are referred to as resource-based policies. CloudFront does not support resource-based policies, but it does support resource-level policies. For more information about the types of permissions policies that CloudFront supports, see [Managing Access to Resources \(p. 479\)](#).

Topics

- [Identity-Based Policies \(IAM Policies\) \(p. 479\)](#)
- [Resource-Level Policies \(p. 480\)](#)
- [Tag-Based Policies \(p. 480\)](#)

Identity-Based Policies (IAM Policies)

You can attach policies to IAM identities. For example, you can do the following:

- **Attach a permissions policy to a user or a group in your account** – An account administrator can use a permissions policy that is associated with a particular user to grant permissions for that user to create a web distribution.
- **Attach a permissions policy to a role (grant cross-account permissions)** – You can grant permissions to perform CloudFront actions to a user that was created in another AWS account. To do so, you attach a permissions policy to an IAM role, and then you allow the user in the other account to assume the role. The following example explains how this works for two AWS accounts, account A and account B:
 1. Account A administrator creates an IAM role and attaches to the role a permissions policy that grants permissions to create or access resources that are owned by account A.
 2. Account A administrator attaches a trust policy to the role. The trust policy identifies account B as the principal that can assume the role.
 3. Account B administrator can then delegate permissions to assume the role to users or groups in account B. This allows users in account B to create or access resources in account A.

For more information about how to delegate permissions to users in another AWS account, see [Access Management](#) in the *IAM User Guide*.

The following example policy allows a user to perform the `CreateDistribution` action to programmatically create a web distribution for your AWS account:

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



```
{
  "Effect": "Allow",
  "Action": [
    "cloudfront:CreateDistribution"
  ],
  "Resource": "*"
}
```

For information about the permissions required to perform operations by using the CloudFront console, see [Permissions Required to Use the CloudFront Console \(p. 483\)](#). For more information about attaching policies to identities for CloudFront, see [Using Identity-Based Policies \(IAM Policies\) for CloudFront \(p. 483\)](#). For more information about users, groups, roles, and permissions, see [Identities \(Users, Groups, and Roles\)](#) in the *IAM User Guide*.

Resource-Level Policies

In some scenarios, you might want to grant a specific level of access to a resource that you specify; for example, access to only specific actions on that resource. One way that some AWS services implement this is to allow you to directly attach a policy on the resource. For example, that's how Amazon S3 and Elasticsearch implement resource access control. CloudFront allows the same flexibility but uses a different method. Instead of attaching a policy to a resource, you specify the resource in a policy.

For example, the following policy shows how you might allow update, delete, and create invalidations access to a distribution that you specify by the distribution's ARN. This policy grants the permissions necessary to complete these actions from the AWS API or AWS CLI only. (To use this policy, replace the italicized text in the example policy with your own resource information.)

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "cloudfront:CreateDistribution",
        "cloudfront:Get*",
        "cloudfront:List*"
      ],
      "Resource": "*"
    },
    {
      "Sid": "VisualEditor1",
      "Effect": "Allow",
      "Action": [
        "cloudfront:UpdateDistribution",
        "cloudfront:DeleteDistribution",
        "cloudfront:CreateInvalidation"
      ],
      "Resource": "arn:aws:cloudfront::123456789012:distribution/EDFDVBD6EXAMPLE"
    }
  ]
}
```

Tag-Based Policies

When you design IAM policies, you might set granular permissions by granting access to specific resources. As the number of resources that you manage grows, this task becomes more difficult. Tagging

resources and using tags in policy statement conditions can make this task easier. You grant access in bulk to any resource with a certain tag. Then you repeatedly apply this tag to relevant resources, during creation or later.

Note

Using tags in conditions is one way to control access to resources and requests. For information about tagging in CloudFront, see [Tagging Amazon CloudFront Distributions \(p. 64\)](#).

Tags can be attached to the resource or passed in the request to services that support tagging. In CloudFront, resources can have tags, and some actions can include tags. When you create an IAM policy, you can use tag condition keys to control:

- Which users can perform actions on a distribution, based on tags that it already has.
- What tags can be passed in an action's request.
- Whether specific tag keys can be used in a request.

For the complete syntax and semantics of tag condition keys, see [Control Access Using IAM Tags](#) in the *IAM User Guide*.

For example, the `AWSCloudFrontFullAccess` managed user policy gives users unrestricted permission to perform any CloudFront action on any resource. The following policy restricts this power and denies unauthorized users permission to create CloudFront production distributions.

To implement the restriction using tags, it denies the `CreateDistribution` action if the request specifies a tag named `stage` with one of the values `gamma` or `prod`. In addition, the policy prevents these unauthorized users from tampering with the stage of production environments by not allowing tag modification actions to include these same tag values or to completely remove the stage tag. A customer's administrator must attach this IAM policy to unauthorized IAM users, in addition to the managed user policy.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": [
        "cloudfront:CreateDistribution",
        "cloudfront:TagResource"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:RequestTag/stage": [
            "gamma",
            "prod"
          ]
        }
      }
    },
    {
      "Effect": "Deny",
      "Action": [
        "cloudfront:TagResource",
        "cloudfront:UntagResource"
      ],
      "Resource": "*",
      "Condition": {
        "ForAnyValue:StringEquals": {
          "aws:TagKeys": [
            "stage"
          ]
        }
      }
    }
  ]
}
```

```
    ],  
    },  
    "StringEquals": {  
      "aws:ResourceTag/stage": [  
        "gamma",  
        "prod"  
      ]  
    }  
  }  
}  
]  
}
```

Specifying Policy Elements: Resources, Actions, Effects, and Principals

CloudFront includes API actions (see [Amazon CloudFront API Reference](#)) that you can use on each CloudFront resource (see [ARNs for CloudFront Resources \(p. 478\)](#)). You can grant a user or a federated user permission to perform any or all of these actions.

The following are the basic policy elements:

- **Resource** – You use an Amazon Resource Name (ARN) to identify the resource that the policy applies to. For more information, see [ARNs for CloudFront Resources \(p. 478\)](#).
- **Action** – You use action keywords to identify resource operations that you want to allow or deny. For example, depending on the specified `Effect`, the `cloudfront:CreateDistribution` permission allows or denies the user permissions to perform the CloudFront `CreateDistribution` action.
- **Effect** – You specify the effect, either allow or deny, when a user tries to perform the action on the specified resource. If you don't explicitly grant access to an action, access is implicitly denied. You can also explicitly deny access to a resource, which you might do to make sure that a user cannot access it, even if a different policy grants access.
- **Principal** – In identity-based policies (IAM policies), the user that the policy is attached to is the implicit principal. For resource-based policies, you specify the user, account, service, or other entity that you want to receive permissions (applies to resource-based policies only).

CloudFront does not support resource-based policies, but it does support resource-level policies. For more information about the types of permissions policies that CloudFront supports, see [Managing Access to Resources \(p. 479\)](#).

For more information about IAM policy syntax and descriptions, see the [AWS IAM Policy Reference](#) in the *IAM User Guide*.

For a list showing all of the CloudFront API operations and the resources that they apply to, see [CloudFront API Permissions: Actions, Resources, and Conditions Reference \(p. 487\)](#).

Specifying Conditions in a Policy

When you grant permissions, you can use the IAM policy language to specify when a policy should take effect. For example, you might want a policy to be applied only after a specific date. For more information about specifying conditions in a policy language, see [Condition](#) in the *IAM User Guide*.

To express conditions, you use predefined condition keys. There are no condition keys specific to CloudFront. However, there are AWS-wide condition keys that you can use as appropriate. For a complete list of AWS-wide keys, see [Available Keys for Conditions](#) in the *IAM User Guide*.

Using Identity-Based Policies (IAM Policies) for CloudFront

This topic provides examples of identity-based policies that demonstrate how an account administrator can attach permissions policies to IAM identities (that is, users, groups, and roles) and thereby grant permissions to perform operations on CloudFront resources.

Important

We recommend that you first review the introductory topics that explain the basic concepts and options to manage access to your CloudFront resources. For more information, see [Overview of Managing Access Permissions to Your CloudFront Resources](#) (p. 478).

Topics

- [Permissions Required to Use the CloudFront Console](#) (p. 483)
- [AWS Managed \(Predefined\) Policies for CloudFront](#) (p. 485)
- [Customer Managed Policy Examples](#) (p. 485)

The following shows a permissions policy. The Sid, or statement ID, is optional.

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

The policy grants permissions to perform all CloudFront operations, which is sufficient to access CloudFront programmatically. If you're using the console to access CloudFront, see [Permissions Required to Use the CloudFront Console](#) (p. 483).

For a list of actions and the ARN that you specify to grant or deny permission to use each action, see [CloudFront API Permissions: Actions, Resources, and Conditions Reference](#) (p. 487).

Permissions Required to Use the CloudFront Console

To grant full access to the CloudFront console, you grant the permissions in the following permissions policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "acm:ListCertificates",
        "cloudfront:*",
        "cloudwatch:DescribeAlarms",
        "cloudwatch:PutMetricAlarm",
        "cloudwatch:GetMetricStatistics",
        "elasticloadbalancing:DescribeLoadBalancers",
        "iam:ListServerCertificates",
        "sns:ListSubscriptionsByTopic",

```

```
        "sns:ListTopics",
        "waf:GetWebACL",
        "waf:ListWebACLs"
    ],
    "Resource": "*"
  },
  {
    "Effect": "Allow",
    "Action": [
      "s3:ListAllMyBuckets",
      "s3:PutBucketPolicy"
    ],
    "Resource": "arn:aws:s3:::*"
  }
]
```

Here's why the permissions are required:

acm:ListCertificates

When you're creating and updating web distributions by using the CloudFront console and you want to configure CloudFront to require HTTPS between the viewer and CloudFront or between CloudFront and the origin, lets you view a list of ACM certificates.

This permission isn't required if you aren't using the CloudFront console.

cloudfront:*

Lets you perform all CloudFront actions.

cloudwatch:DescribeAlarms and cloudwatch:PutMetricAlarm

Let you create and view CloudWatch alarms in the CloudFront console. See also `sns:ListSubscriptionsByTopic` and `sns:ListTopics`.

These permissions aren't required if you aren't using the CloudFront console.

cloudwatch:GetMetricStatistics

Lets CloudFront render CloudWatch metrics in the CloudFront console.

This permission isn't required if you aren't using the CloudFront console.

elasticloadbalancing:DescribeLoadBalancers

When creating and updating web distributions, lets you view a list of Elastic Load Balancing load balancers in the list of available origins.

This permission isn't required if you aren't using the CloudFront console.

iam:ListServerCertificates

When you're creating and updating web distributions by using the CloudFront console and you want to configure CloudFront to require HTTPS between the viewer and CloudFront or between CloudFront and the origin, lets you view a list of certificates in the IAM certificate store.

This permission isn't required if you aren't using the CloudFront console.

s3:ListAllMyBuckets

When you're creating and updating web and RTMP distributions, lets you perform the following operations:

- View a list of S3 buckets in the list of available origins

- View a list of S3 buckets that you can save access logs in

This permission isn't required if you aren't using the CloudFront console.

S3:PutBucketPolicy

When you're creating or updating distributions that restrict access to S3 buckets, lets a user update the bucket policy to grant access to the CloudFront origin access identity. For more information, see [Restricting Access to Amazon S3 Content by Using an Origin Access Identity \(p. 209\)](#).

This permission isn't required if you aren't using the CloudFront console.

sns:ListSubscriptionsByTopic and sns:ListTopics

When you create CloudWatch alarms in the CloudFront console, lets you choose an SNS topic for notifications.

These permissions aren't required if you aren't using the CloudFront console.

waf:GetWebACL and waf:ListWebACLs

Lets you view a list of AWS WAF web ACLs in the CloudFront console.

These permissions aren't required if you aren't using the CloudFront console.

AWS Managed (Predefined) Policies for CloudFront

AWS addresses many common use cases by providing standalone IAM policies that are created and administered by AWS. These AWS managed policies grant necessary permissions for common use cases so that you can avoid having to investigate what permissions are needed. For more information, see [AWS Managed Policies](#) in the *IAM User Guide*. For CloudFront, IAM provides two managed policies:

- **CloudFrontFullAccess** – Grants full access to CloudFront resources.

Important

If you want CloudFront to create and save access logs, you need to grant additional permissions. For more information, see [Permissions Required to Configure Standard Logging and to Access Your Log Files \(p. 440\)](#).

- **CloudFrontReadOnlyAccess** – Grants read-only access to CloudFront resources.

Note

You can review these permissions policies by signing in to the IAM console and searching for specific policies there. You can also create your own custom IAM policies to allow permissions for CloudFront API operations. You can attach these custom policies to the IAM users or groups that require those permissions.

Customer Managed Policy Examples

You can create your own custom IAM policies to allow permissions for CloudFront API actions. You can attach these custom policies to the IAM users or groups that require the specified permissions. These policies work when you are using the CloudFront API, the AWS SDKs, or the AWS CLI. The following examples show permissions for a few common use cases. For the policy that grants a user full access to CloudFront, see [Permissions Required to Use the CloudFront Console \(p. 483\)](#).

Examples

- [Example 1: Allow Read Access to All Web Distributions \(p. 486\)](#)
- [Example 2: Allow Creation, Updating, and Deletion of Web Distributions \(p. 486\)](#)
- [Example 3: Allow Creation and Listing of Invalidations \(p. 487\)](#)

Example 1: Allow Read Access to All Web Distributions

The following permissions policy grants the user permissions to view all web distributions in the CloudFront console:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "acm:ListCertificates",
        "cloudfront:GetDistribution",
        "cloudfront:GetDistributionConfig",
        "cloudfront:ListDistributions",
        "cloudfront:ListCloudFrontOriginAccessIdentities",
        "elasticloadbalancing:DescribeLoadBalancers",
        "iam:ListServerCertificates",
        "sns:ListSubscriptionsByTopic",
        "sns:ListTopics",
        "waf:GetWebACL",
        "waf:ListWebACLs"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "s3:ListAllMyBuckets"
      ],
      "Resource": "arn:aws:s3:::*"
    }
  ]
}
```

Example 2: Allow Creation, Updating, and Deletion of Web Distributions

The following permissions policy allows users to create, update, and delete web distributions by using the CloudFront console:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "acm:ListCertificates",
        "cloudfront:CreateDistribution",
        "cloudfront>DeleteDistribution",
        "cloudfront:GetDistribution",
        "cloudfront:GetDistributionConfig",
        "cloudfront:ListDistributions",
        "cloudfront:UpdateDistribution",
        "cloudfront:ListCloudFrontOriginAccessIdentities",
        "elasticloadbalancing:DescribeLoadBalancers",
        "iam:ListServerCertificates",
        "sns:ListSubscriptionsByTopic",
        "sns:ListTopics",
        "waf:GetWebACL",
        "waf:ListWebACLs"
      ],
      "Resource": "*"
    }
  ],
}
```

```
{
  "Effect": "Allow",
  "Action": [
    "s3:ListAllMyBuckets",
    "s3:PutBucketPolicy"
  ],
  "Resource": "arn:aws:s3:::*"
}
```

The `cloudfront:ListCloudFrontOriginAccessIdentities` permission allows users to automatically grant to an existing origin access identity the permission to access objects in an Amazon S3 bucket. If you also want users to be able to create origin access identities, you also need to allow the `cloudfront:CreateCloudFrontOriginAccessIdentity` permission.

Example 3: Allow Creation and Listing of Invalidations

The following permissions policy allows users to create and list invalidations. It includes read access to CloudFront distributions because you create and view invalidations by first displaying settings for a distribution:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "acm:ListCertificates",
        "cloudfront:GetDistribution",
        "cloudfront:GetStreamingDistribution",
        "cloudfront:GetDistributionConfig",
        "cloudfront:ListDistributions",
        "cloudfront:ListCloudFrontOriginAccessIdentities",
        "cloudfront:CreateInvalidation",
        "cloudfront:GetInvalidation",
        "cloudfront:ListInvalidations",
        "elasticloadbalancing:DescribeLoadBalancers",
        "iam:ListServerCertificates",
        "sns:ListSubscriptionsByTopic",
        "sns:ListTopics",
        "waf:GetWebACL",
        "waf:ListWebACLs"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "s3:ListAllMyBuckets"
      ],
      "Resource": "arn:aws:s3:::*"
    }
  ]
}
```

CloudFront API Permissions: Actions, Resources, and Conditions Reference

When you are setting up [Access Control \(p. 478\)](#) and writing a permissions policy that you can attach to an IAM identity (identity-based policies), you can use the following lists as a reference. The lists

include each CloudFront API operation, the corresponding actions for which you can grant permissions to perform the action, and the AWS resource for which you can grant the permissions. You specify the actions in the policy's `Action` field, and you specify the resource value in the policy's `Resource` field.

You can use AWS-wide condition keys in your CloudFront policies to express conditions. For a complete list of AWS-wide keys, see [Available Keys](#) in the *IAM User Guide*.

Topics

- [Required Permissions for Actions on Web Distributions](#) (p. 488)
- [Required Permissions for Actions on RTMP Distributions](#) (p. 489)
- [Required Permissions for Actions on Invalidations](#) (p. 490)
- [Required Permissions for Actions on Origin Access Identities](#) (p. 491)
- [Required Permissions for CloudFront Actions Related to Lambda@Edge](#) (p. 491)
- [Required Permissions for Actions on Tags](#) (p. 492)

Required Permissions for Actions on Web Distributions

CreateDistribution

Required Permissions (API Action):

- `cloudfront:CreateDistribution`
- `acm:ListCertificates` (CloudFront console only)
- Only if you configure CloudFront to save access logs:
 - `s3:GetBucketAcl`
 - `s3:PutBucketAcl`
 - The S3 ACL for the bucket must grant you `FULL_CONTROL`

Resources:

- CloudFront: *
- ACM: *
- Amazon S3: If you configure CloudFront to save access logs, you can optionally restrict access to a specified bucket.

CreateDistributionWithTags

Required Permissions (API Action):

- `cloudfront:CreateDistribution`, `cloudfront:TagResource`
- `acm:ListCertificates` (CloudFront console only)
- Only if you configure CloudFront to save access logs:
 - `s3:GetBucketAcl`
 - `s3:PutBucketAcl`
 - The S3 ACL for the bucket must grant you `FULL_CONTROL`

Resources:

- CloudFront: *
- ACM: *
- Amazon S3: If you configure CloudFront to save access logs, you can optionally restrict access to a specified bucket.

GetDistribution

Required Permissions (API Action): `cloudfront:GetDistribution`, `acm:ListCertificates` (CloudFront console only)

Resources: *

[GetDistributionConfig](#)

Required Permissions (API Action): `cloudfront:GetDistributionConfig`,
`acm:ListCertificates` (CloudFront console only)

Resources: *

[ListDistributions](#)

Required Permissions (API Action): `cloudfront:ListDistributions`

Resources: *

[UpdateDistribution](#)

Required Permissions (API Action):

- `cloudfront:UpdateDistribution`
- `acm:ListCertificates` (CloudFront console only)
- Only if you configure CloudFront to save access logs:
 - `s3:GetBucketAcl`
 - `s3:PutBucketAcl`
 - The S3 ACL for the bucket must grant you `FULL_CONTROL`

Resources:

- CloudFront: *
- ACM: *
- Amazon S3: If you configure CloudFront to save access logs, you can optionally restrict access to a specified bucket.

[DeleteDistribution](#)

Required Permissions (API Action): `cloudfront:DeleteDistribution`

Resources: *

Required Permissions for Actions on RTMP Distributions

[CreateStreamingDistribution](#)

Required Permissions (API Action): `cloudfront:CreateStreamingDistribution`

Only if you configure CloudFront to save access logs:

- `s3:GetBucketAcl`
- `s3:PutBucketAcl`
- The S3 ACL for the bucket must grant you `FULL_CONTROL`

Resources: *

If you configure CloudFront to save access logs, you can optionally restrict access to a specified bucket.

[CreateStreamingDistributionWithTags](#)

Required Permissions (API Action): `cloudfront:CreateStreamingDistribution`,
`cloudfront:TagResource`

Only if you configure CloudFront to save access logs:

- `s3:GetBucketAcl`
- `s3:PutBucketAcl`
- The S3 ACL for the bucket must grant you `FULL_CONTROL`

Resources: *

If you configure CloudFront to save access logs, you can optionally restrict access to a specified bucket.

[GetStreamingDistribution](#)

Required Permissions (API Action): `cloudfront:GetStreamingDistribution`

Resources: *

[GetStreamingDistributionConfig](#)

Required Permissions (API Action): `cloudfront:GetStreamingDistributionConfig`

Resources: *

[ListStreamingDistributions](#)

Required Permissions (API Action): `cloudfront:ListStreamingDistributions`

Resources: *

[UpdateStreamingDistribution](#)

Required Permissions (API Action): `cloudfront:UpdateStreamingDistribution`

Only if you configure CloudFront to save access logs:

- `s3:GetBucketAcl`
- `s3:PutBucketAcl`
- The S3 ACL for the bucket must grant you `FULL_CONTROL`

Resources: *

If you configure CloudFront to save access logs, you can optionally restrict access to a specified bucket.

[DeleteStreamingDistribution](#)

Required Permissions (API Action): `cloudfront:DeleteDistribution`

Resources: *

Required Permissions for Actions on Invalidations

[CreateInvalidation](#)

Required Permissions (API Action): `cloudfront:CreateInvalidation`

Resources: *

[GetInvalidation](#)

Required Permissions (API Action): `cloudfront:GetInvalidation`

Resources: *

ListInvalidations

Required Permissions (API Action): `cloudfront:ListInvalidations`

Resources: *

Required Permissions for Actions on Origin Access Identities

CreateCloudFrontOriginAccessIdentity

Required Permissions (API Action): `cloudfront:CreateCloudFrontOriginAccessIdentity`

Resources: *

GetCloudFrontOriginAccessIdentity

Required Permissions (API Action): `cloudfront:GetCloudFrontOriginAccessIdentity`

Resources: *

GetCloudFrontOriginAccessIdentityConfig

Required Permissions (API Action):
`cloudfront:GetCloudFrontOriginAccessIdentityConfig`

Resources: *

ListCloudFrontOriginAccessIdentities

Required Permissions (API Action): `cloudfront:ListCloudFrontOriginAccessIdentities`

Resources: *

UpdateCloudFrontOriginAccessIdentity

Required Permissions (API Action): `cloudfront:UpdateCloudFrontOriginAccessIdentity`

Resources: *

DeleteCloudFrontOriginAccessIdentity

Required Permissions (API Action): `cloudfront>DeleteCloudFrontOriginAccessIdentity`

Resources: *

Required Permissions for CloudFront Actions Related to Lambda@Edge

To use Lambda@Edge, you need the following CloudFront permissions so you can create or update a distribution that includes triggers for Lambda functions. For information about the Lambda permissions that you need, see [Setting IAM Permissions](#) in the "AWS Lambda@Edge" chapter in the *AWS Lambda Developer Guide*.

CreateDistribution

Required Permissions (API Action):

- `cloudfront:CreateDistribution`
- `acm:ListCertificates` (CloudFront console only)
- Only if you configure CloudFront to save access logs:
 - `s3:GetBucketAcl`

- `s3:PutBucketAcl`
- The S3 ACL for the bucket must grant you `FULL_CONTROL`

Resources:

- CloudFront: *
- ACM: *
- Amazon S3: If you configure CloudFront to save access logs, you can optionally restrict access to a specified bucket.

[CreateDistributionWithTags](#)

Required Permissions (API Action):

- `cloudfront:CreateDistribution`, `cloudfront:TagResource`
- `acm:ListCertificates` (CloudFront console only)
- Only if you configure CloudFront to save access logs:
 - `s3:GetBucketAcl`
 - `s3:PutBucketAcl`
 - The S3 ACL for the bucket must grant you `FULL_CONTROL`

Resources:

- CloudFront: *
- ACM: *
- Amazon S3: If you configure CloudFront to save access logs, you can optionally restrict access to a specified bucket.

[UpdateDistribution](#)

Required Permissions (API Action):

- `cloudfront:UpdateDistribution`
- `acm:ListCertificates` (CloudFront console only)
- Only if you configure CloudFront to save access logs:
 - `s3:GetBucketAcl`
 - `s3:PutBucketAcl`
 - The S3 ACL for the bucket must grant you `FULL_CONTROL`

Resources:

- CloudFront: *
- ACM: *
- Amazon S3: If you configure CloudFront to save access logs, you can optionally restrict access to a specified bucket.

Required Permissions for Actions on Tags

[TagResource](#)

Required Permissions (API Action): `cloudfront:TagResource`

Resources: *

[UntagResource](#)

Required Permissions (API Action): `cloudfront:UntagResource`

Resources: *

ListTagsForResource

Required Permissions (API Action): `cloudfront:ListTagsForResource`

Resources: *

Logging and Monitoring in Amazon CloudFront

Monitoring is an important part of maintaining the availability and performance of CloudFront and your AWS solutions. You should collect monitoring data from all of the parts of your AWS solution so that you can more easily debug a multi-point failure if one occurs. AWS provides several tools for monitoring your CloudFront resources and activity, and responding to potential incidents:

Amazon CloudWatch Alarms

Using CloudWatch alarms, you watch a single metric over a time period that you specify. If the metric exceeds a given threshold, a notification is sent to an Amazon SNS topic or AWS Auto Scaling policy. CloudWatch alarms do not invoke actions when a metric is in a particular state. Rather the state must have changed and been maintained for a specified number of periods. For more information, see [Monitoring CloudFront with Amazon CloudWatch \(p. 428\)](#).

AWS CloudTrail Logs

CloudTrail provides a record of actions taken by a user, role, or an AWS service in CloudFront. Using the information collected by CloudTrail, you can determine the request that was made to CloudFront, the IP address from which the request was made, who made the request, when it was made, and additional details. For more information, see [Using AWS CloudTrail to Capture Requests Sent to the CloudFront API \(p. 466\)](#).

CloudFront Access Logs

Server access logs provide detailed records about requests that are made to a distribution. Server access logs are useful for many applications. For example, access log information can be useful in security and access audits. For more information, see [Configuring and Using Standard Logs \(Access Logs\) \(p. 439\)](#).

CloudFront Console Reports

The CloudFront console includes a variety of reports, including the cache statistics report, the popular objects report, and the top referrers report. Most CloudFront console reports are based on the data in CloudFront access logs, which contain detailed information about every user request that CloudFront receives. However, you don't need to enable access logs to view the reports. For more information, see [CloudFront Reports in the Console \(p. 405\)](#).

Compliance Validation for Amazon CloudFront

Third-party auditors assess the security and compliance of Amazon CloudFront as part of multiple AWS compliance programs. These include SOC, PCI, HIPAA, and others.

For a list of AWS services in scope of specific compliance programs, see [AWS Services in Scope by Compliance Program](#). For general information, see [AWS Compliance Programs](#).

You can download third-party audit reports using AWS Artifact. For more information, see [Downloading Reports in AWS Artifact](#).

Your compliance responsibility when using CloudFront is determined by the sensitivity of your data, your company's compliance objectives, and applicable laws and regulations. AWS provides the following resources to help with compliance:

- [Security and Compliance Quick Start Guides](#) – These deployment guides discuss architectural considerations and provide steps for deploying security- and compliance-focused baseline environments on AWS.
- [Architecting for HIPAA Security and Compliance Whitepaper](#) – This whitepaper describes how companies can use AWS to create HIPAA-compliant applications.

The AWS HIPAA compliance program includes CloudFront as a HIPAA eligible service. If you have an executed Business Associate Addendum (BAA) with AWS, you can use CloudFront to deliver content that contains protected health information (PHI). For more information, see [HIPAA Compliance](#).

- [AWS Compliance Resources](#) – This collection of workbooks and guides might apply to your industry and location.
- [AWS Config](#) – This AWS service assesses how well your resource configurations comply with internal practices, industry guidelines, and regulations.
- [AWS Security Hub](#) – This AWS service provides a comprehensive view of your security state within AWS that helps you check your compliance with security industry standards and best practices.

CloudFront Compliance Best Practices

This section provides best practices and recommendations for compliance when you use Amazon CloudFront to serve your content.

If you run PCI-compliant or HIPAA-compliant workloads that are based on the [AWS shared responsibility model](#), we recommend that you log your CloudFront usage data for the last 365 days for future auditing purposes. To log usage data, you can do the following:

- Enable CloudFront access logs. For more information, see [Configuring and Using Standard Logs \(Access Logs\)](#) (p. 439).
- Capture requests that are sent to the CloudFront API. For more information, see [Using AWS CloudTrail to Capture Requests Sent to the CloudFront API](#) (p. 466).

In addition, see the following for details about how CloudFront is compliant with the PCI DSS and SOC standards.

PCI DSS

CloudFront supports the processing, storage, and transmission of credit card data by a merchant or service provider, and has been validated as being compliant with Payment Card Industry (PCI) Data Security Standard (DSS). For more information about PCI DSS, including how to request a copy of the AWS PCI Compliance Package, see [PCI DSS Level 1](#).

As a security best practice, we recommend that you don't cache credit card information in CloudFront edge caches. For example, you can configure your origin to include a `Cache-Control:no-cache="field-name"` header in responses that contain credit card information, such as the last four digits of a credit card number and the card owner's contact information.

AWS System and Organization Controls

CloudFront is compliant with System and Organization Controls (SOC) measures, including SOC 1, SOC 2, and SOC 3. SOC reports are independent, third-party examination reports that demonstrate how AWS achieves key compliance controls and objectives. These audits ensure that the appropriate safeguards and procedures are in place to protect against risks that might affect the security, confidentiality, and availability of customer and company data. The results of these third-party audits are available on the [AWS SOC Compliance website](#), where you can view the published reports to get more information about the controls that support AWS operations and compliance.

Resilience in Amazon CloudFront

The AWS global infrastructure is built around AWS Regions and Availability Zones. AWS Regions provide multiple physically separated and isolated Availability Zones, which are connected with low-latency, high-throughput, and highly redundant networking. With Availability Zones, you can design and operate applications and databases that automatically fail over between Availability Zones without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures.

For more information about AWS Regions and Availability Zones, see [AWS Global Infrastructure](#).

CloudFront Origin Failover

In addition to the support of AWS global infrastructure, Amazon CloudFront offers an *origin failover* feature to help support your data resiliency needs. CloudFront is a global service that delivers your content through a worldwide network of data centers called *edge locations* or *points of presence* (POPs). If your content is not already cached in an edge location, CloudFront retrieves it from an origin that you've identified as the source for the definitive version of the content.

You can improve resiliency and increase availability for specific scenarios by setting up CloudFront with origin failover. To get started, you create an origin group in which you designate a primary origin for CloudFront plus a second origin. CloudFront automatically switches to the second origin when the primary origin returns specific HTTP status code failure responses. For more information, see [Optimizing High Availability with CloudFront Origin Failover](#) (p. 231).

Infrastructure Security in Amazon CloudFront

As a managed service, Amazon CloudFront is protected by the AWS global network security procedures that are described in the [Amazon Web Services: Overview of Security Processes](#) whitepaper.

You use AWS published API calls to access CloudFront through the network. Clients must support Transport Layer Security (TLS) 1.0 or later. We recommend TLS 1.2 or later. Clients must also support cipher suites with perfect forward secrecy (PFS) such as Ephemeral Diffie-Hellman (DHE) or Elliptic Curve Ephemeral Diffie-Hellman (ECDHE). Most modern systems such as Java 7 and later support these modes. Additionally, requests must be signed by using an access key ID and a secret access key that is associated with an IAM principal. Or you can use the [AWS Security Token Service](#) (AWS STS) to generate temporary security credentials to sign requests.

Quotas

CloudFront is subject to the following quotas (formerly referred to as limits). Note that Lambda@Edge has specific quotas as well, that are in addition to the default CloudFront quotas.

Topics

- [General Quotas \(p. 496\)](#)
- [General Quotas on Web Distributions \(p. 497\)](#)
- [General Quotas on Policies \(p. 497\)](#)
- [Quotas on WebSocket Connections \(p. 498\)](#)
- [Quotas on Whitelisted Cookies \(Web Distributions Only\) \(p. 498\)](#)
- [Quotas on Whitelisted Query Strings \(Web Distributions Only\) \(p. 499\)](#)
- [Quotas on Custom Headers \(Web Distributions Only\) \(p. 499\)](#)
- [Quotas on SSL Certificates \(Web Distributions Only\) \(p. 499\)](#)
- [Quotas on Invalidations \(p. 500\)](#)
- [Quotas on Field-Level Encryption \(p. 500\)](#)
- [Quotas on Lambda@Edge \(p. 500\)](#)
- [Request Timeout \(p. 502\)](#)
- [Quotas on RTMP Distributions \(p. 502\)](#)

General Quotas

Entity	Default quota
Data transfer rate per distribution	150 Gbps Request a higher quota
Requests per second per distribution	250,000 Request a higher quota
Tags that can be added to a distribution	50
Files that you can serve per distribution	No quota
Maximum length of a request, including headers and query strings, but not including the body content	20,480 bytes
Maximum length of a URL	8,192 bytes
Active CloudFront key pairs for trusted signers	2
For more information, see Specifying the AWS Accounts That Can Create Signed URLs and Signed Cookies (Trusted Signers) (p. 150) .	

General Quotas on Web Distributions

Entity	Default quota
Web distributions per AWS account For more information, see Creating a Distribution (p. 37) .	200 Request a higher quota
Maximum file size for HTTP GET, POST, and PUT requests	20 GB
Response timeout per origin For more information, see Origin Response Timeout (p. 45) .	1-60 seconds Request a higher quota
Connection timeout per origin For more information, see Origin Connection Timeout (p. 43) .	1-10 seconds
Connection attempts per origin For more information, see Origin Connection Attempts (p. 42) .	1-3
File compression: range of file sizes that CloudFront compresses For more information, see Serving compressed files (p. 116) .	1,000 to 10,000,000 bytes
Alternate domain names (CNAMEs) per distribution For more information, see Using Custom URLs for Files by Adding Alternate Domain Names (CNAMEs) (p. 71) .	100 Request a higher quota
Origins per distribution	25 Request a higher quota
Origin groups per distribution	10 Request a higher quota
Origin access identities per account	100 Request a higher quota
Cache behaviors per distribution	25 Request a higher quota

General Quotas on Policies

Entity	Default quota
Cache policies per AWS account	20
Distributions associated with the same cache policy	100
Query strings per cache policy	10

Entity	Default quota
	Request a higher quota
Headers per cache policy	10 Request a higher quota
Cookies per cache policy	10 Request a higher quota
Origin request policies per AWS account	20
Distributions associated with the same origin request policy	100
Query strings per origin request policy	10 Request a higher quota
Headers per origin request policy	10 Request a higher quota
Cookies per origin request policy	10 Request a higher quota

Quotas on WebSocket Connections

Entity	Default quota
Origin response timeout (idle timeout)	10 minutes If CloudFront hasn't detected any bytes sent from the origin to the client within the past 10 minutes, the connection is assumed to be idle and is closed.

Quotas on Whitelisted Cookies (Web Distributions Only)

Entity	Default quota
Whitelisted cookies per cache behavior	10
For more information, see Caching Content Based on Cookies (p. 244) .	Request a higher quota
Total number of bytes in whitelisted cookie names (doesn't apply if you configure CloudFront to forward all cookies to the origin)	512 minus the number of whitelisted cookies

Quotas on Whitelisted Query Strings (Web Distributions Only)

Entity	Default quota
Maximum number of characters in a whitelisted query string	128 characters
Maximum number of characters total for all whitelisted query strings in the same parameter	512 characters
Whitelisted query strings per cache behavior	10
For more information, see Caching Content Based on Query String Parameters (p. 241).	Request a higher quota

Quotas on Custom Headers (Web Distributions Only)

Entity	Default quota
Whitelisted headers per cache behavior	10
For more information, see Caching Content Based on Request Headers (p. 246).	Request a higher quota
Custom headers: maximum number of custom headers that you can configure CloudFront to add to origin requests	10 name/value pairs
For more information, see Adding Custom Headers to Origin Requests (p. 282).	Request a higher quota
Custom headers: maximum length of a header name	256 characters
Custom headers: maximum length of a header value	1,783 characters
Custom headers: maximum length of all header values and names combined	10,240 characters

Quotas on SSL Certificates (Web Distributions Only)

Entity	Default quota
SSL certificates per AWS account when serving HTTPS requests using dedicated IP addresses (no quota when serving HTTPS requests using SNI)	2
For more information, see Using HTTPS with CloudFront (p. 121).	Request a higher quota
SSL certificates that can be associated with a CloudFront web distribution	1

Quotas on Invalidations

Entity	Default quota
File invalidation: maximum number of files allowed in active invalidation requests, excluding wildcard invalidations For more information, see Invalidating Files (p. 110) .	3,000
File invalidation: maximum number of active wildcard invalidations allowed	15
File invalidation: maximum number of files that one wildcard invalidation can process	No quota

Quotas on Field-Level Encryption

Entity	Default quota
Maximum length of a field to encrypt For more information, see Using Field-Level Encryption to Help Protect Sensitive Data (p. 219) .	16 KB
Maximum number of fields in a request body when field-level encryption is configured	10
Maximum length of a request body when field-level encryption is configured	1 MB
Maximum number of field-level encryption configurations that can be associated with one AWS account	10
Maximum number of field-level encryption profiles that can be associated with one AWS account	10
Maximum number of public keys that can be added to one AWS account	10
Maximum number of fields to encrypt that can be specified in one profile	10
Maximum number of CloudFront distributions that can be associated with a field-level encryption configuration	20
Maximum number of query argument profile mappings that can be included in a field-level encryption configuration	5

Quotas on Lambda@Edge

The quotas in this section apply to Lambda@Edge. These quotas are in addition to the default CloudFront and Lambda quotas, which also apply. For the default quotas, see [Quotas \(p. 496\)](#) in this guide and [Quotas](#) in the *AWS Lambda Developer Guide*.

Note

Lambda dynamically scales capacity in response to increased traffic, within your account's quotas. For more information, see [Function Scaling](#) in the *AWS Lambda Developer Guide*.

In addition, be aware that there are some other restrictions when using Lambda@Edge functions. For more information, see [Requirements and Restrictions on Lambda Functions \(p. 393\)](#).

Quotas that differ by event type

Entity	Origin request and response event quotas	Viewer request and response event quotas
Function memory size	Same as Lambda quotas	128 MB
Function timeout. The function can make network calls to resources such as Amazon S3 buckets, DynamoDB tables, or Amazon EC2 instances in AWS Regions.	30 seconds	5 seconds
Size of a response that is generated by a Lambda function, including headers and body	1 MB	40 KB
Maximum compressed size of a Lambda function and any included libraries	50 MB	1 MB

Other quotas

Entity	Default quota
Distributions per AWS account that you can create triggers for	25 Request a higher quota
Triggers per distribution	100 Request a higher quota
Requests per second	10,000 (in each Region) Request a higher quota
Concurrent executions	1000 (in each Region)
For more information, see Function Scaling in the <i>AWS Lambda Developer Guide</i> .	Request a higher quota

Size Quotas on URI and Query String

When accessing or updating a URI or query string in a Lambda@Edge function, the total length of the URI including the query string must be less than 8,192 characters.

Size Quotas on Request Body with the Include Body Option

When you choose the **Include Body** option to expose the request body to your Lambda@Edge function, the following size quotas apply to the portions of the body that are exposed or replaced.

Note

The body is always base64 encoded by Lambda@Edge before it is exposed.

Size Quotas when Exposing the Body to a Lambda Function

If the request body is large, Lambda@Edge truncates it before exposing it, as follows:

- For viewer requests, the body is truncated at 40 KB.
- For origin requests, the body is truncated at 1 MB.

Size Quotas when Returning a Request Body from a Lambda Function

If you access the request body as read-only, the full original request body is sent to the origin. However, if you choose to replace the request body, the following body size quotas apply when it's returned from a Lambda function:

Type of body encoding	Allowed body size: Viewer request	Allowed body size: Origin request
text	40 KB	1 MB
base64	53.2 KB	1.33 MB

Request Timeout

Entity	Default quota
Request timeout	30 seconds
For more information, see Origin Response Timeout (p. 278) .	Request a higher quota

Quotas on RTMP Distributions

Entity	Default quota
RTMP distributions per AWS account	100
For more information, see Working with RTMP Distributions (p. 305) .	Request a higher quota

Amazon CloudFront Related Information

The information and resources listed here can help you learn more about CloudFront.

Topics

- [Additional Amazon CloudFront Documentation](#) (p. 503)
- [Getting Support](#) (p. 503)
- [CloudFront Developer Tools and SDKs](#) (p. 503)
- [Tips from the Amazon Web Services Blog](#) (p. 504)

Additional Amazon CloudFront Documentation

The following related resources can help you as you work with this service.

- [Amazon CloudFront API Reference](#) – Gives complete descriptions of the API actions, parameters, and data types, and a list of errors that the service returns.
- [CloudFront What's New](#) – Announcements of new CloudFront features and recently added edge locations.
- [Amazon CloudFront product information](#) – The primary web page for information about CloudFront, including features and pricing information.

Getting Support

Support for CloudFront is available in a number of forms.

- [Discussion forums](#) – A community-based forum for developers to discuss technical questions related to CloudFront.
- [AWS Support Center](#) – This site brings together information about your recent support cases and results from AWS Trusted Advisor and health checks, as well as providing links to discussion forums, technical FAQs, the service health dashboard, and information about AWS support plans.
- [AWS Premium Support Information](#) – The primary web page for information about AWS Premium Support, a one-on-one, fast-response support channel to help you build and run applications on AWS Infrastructure Services.
- [AWS IQ](#) – Get help from AWS Certified professionals and experts.
- [Contact Us](#) – Links for inquiring about your billing or account. For technical questions, use the discussion forums or support links above.

CloudFront Developer Tools and SDKs

See the [Developer Tools](#) page for links to developer resources that provide documentation, code samples, release notes, and other information to help you build innovative applications with AWS.

In addition, Amazon Web Services provides software development kits for accessing CloudFront programmatically. The SDK libraries automate a number of common tasks, including cryptographically signing your service requests, retrying requests, and handling error responses.

Tips from the Amazon Web Services Blog

The AWS Blog has a number of posts to help you use CloudFront. For example, see the following blog posts about using CloudFront with WordPress: [Accelerating WordPress with CloudFront using the AWS for WordPress Plugin](#).

Document history

The following entries describe important changes made to the CloudFront documentation.

Change	Description	Date changed
New compression format	CloudFront now supports the Brotli compression format when you configure CloudFront to compress objects at CloudFront edge locations. You can also configure CloudFront to cache Brotli objects using a normalized <code>Accept-Encoding</code> header. For more information, see Serving compressed files (p. 116) and Cache compressed objects (uses the Accept-Encoding header) (p. 86) .	September 14, 2020
New TLS protocol	CloudFront now supports the TLS 1.3 protocol for HTTPS connections between viewers and CloudFront distributions. TLS 1.3 is enabled by default in all CloudFront security policies. For more information, see Supported protocols and ciphers between viewers and CloudFront (p. 128) .	September 3, 2020
New real-time logs	CloudFront now supports configurable real-time logs. With real-time logs, you can get information about requests made to a distribution in real time. You can use real-time logs to monitor, analyze, and take action based on content delivery performance. For more information, see Real-time logs (p. 457) .	August 31, 2020
API support for additional metrics	CloudFront now supports enabling eight additional real-time metrics with the CloudFront API. For more information, see Enabling additional metrics (p. 430) .	August 28, 2020
New CloudFront HTTP headers	CloudFront added additional HTTP headers for determining information about the viewer such as device type, geographic location, and more. For more information, see Using the CloudFront HTTP headers (p. 102) .	July 23, 2020
New feature	CloudFront now supports <i>cache policies</i> and <i>origin request policies</i> , which give you more granular control over the <i>cache key</i> and <i>origin requests</i> for your CloudFront distributions. For more information, see Working with policies (p. 83) .	July 22, 2020
New security policy	CloudFront now supports a new security policy, TLSv1.2_2019 , with a smaller set of supported ciphers. For more information, see Supported protocols and ciphers between viewers and CloudFront (p. 128) .	July 8, 2020
New settings to control origin timeouts and attempts	CloudFront added new settings that control origin timeouts and attempts. For more information, see Controlling Origin Timeouts and Attempts (p. 233) .	June 5, 2020
New documentation for getting started with CloudFront by creating a	Get started with CloudFront by creating a secure static website using Amazon S3, CloudFront, Lambda@Edge, and more, all deployed with AWS CloudFormation. For more information, see Getting started with a secure static website (p. 28) .	June 2, 2020

Change	Description	Date changed
secure static website		
Lambda@Edge supports newer runtime versions	Lambda@Edge now supports Lambda functions with the Node.js 12 and Python 3.8 runtimes. For more information, see Lambda Function Supported Runtimes and Configuration (p. 397) .	February 27, 2020
New real-time metrics in CloudWatch	Amazon CloudFront now offers eight additional real-time metrics in Amazon CloudWatch. For more information, see Viewing additional CloudFront distribution metrics (p. 430) .	December 19, 2019
New fields in access logs	CloudFront adds seven new fields to access logs. For more information, see Web Distribution Standard Log File Format (p. 445) .	December 12, 2019
AWS for WordPress plugin	You can use the AWS for WordPress plugin to provide visitors to your WordPress website an accelerated viewing experience using CloudFront. For more information, see Getting started with a simple distribution (p. 19) .	October 30, 2019
Tag-based and resource-level IAM permissions policies	CloudFront now supports two additional ways of specifying IAM permission policies: tag-based and resource-level policy permissions. For more information, see Managing Access to Resources .	August 8, 2019
Support for Python programming language	You can now use the Python programming language to develop functions in Lambda@Edge, in addition to Node.js. For example functions that cover a variety of scenarios, see Lambda@Edge Example Functions .	August 1, 2019
Updated monitoring graphs	Content updates to describe new ways for you to monitor Lambda functions associated with your CloudFront distributions directly from the CloudFront console to more easily track and debug errors. For more information, see Monitoring CloudFront .	June 20, 2019
Consolidated security content	A new Security chapter consolidates information about CloudFront's features around and implementation of data protection, IAM, logging, compliance, and more. For more information, see Security .	May 24, 2019
Domain validation is now required	CloudFront now requires that you use an SSL certificate to verify that you have permission to use an alternate domain name with a distribution. For more information, see Using Alternate Domain Names and HTTPS .	April 9, 2019
Updated PDF filename	The new filename for the Amazon CloudFront Developer Guide is: AmazonCloudFront_DevGuide. The previous name was: cf-dg.	January 7, 2019
New features	CloudFront now supports WebSocket, a TCP-based protocol that is useful when you need long-lived connections between clients and servers. You can also now set up CloudFront with origin failover for scenarios that require high availability. For more information, see Using WebSocket with CloudFront Distributions and Optimizing High Availability with CloudFront Origin Failover .	November 20, 2018

Change	Description	Date changed
New feature	CloudFront now supports detailed error logging for HTTP requests that run Lambda functions. You can store the logs in CloudWatch and use them to help troubleshoot HTTP 5xx errors when your function returns an invalid response. For more information, see CloudWatch Metrics and CloudWatch Logs for Lambda Functions .	October 8, 2018
New feature	You can now opt to have Lambda@Edge expose the body in a request for writable HTTP methods (POST, PUT, DELETE, and so on), so that you can access it in your Lambda function. You can choose read-only access, or you can specify that you'll replace the body. For more information, see Accessing the Request Body by Choosing the Include Body Option .	August 14, 2018
New feature	CloudFront now supports serving content compressed by using brotli or other compression algorithms, in addition to or instead of gzip. For more information, see Serving Compressed Files .	July 25, 2018
Reorganization	The Amazon CloudFront Developer Guide has been reorganized to simplify finding related content, and to improve scanability and navigation.	June 28, 2018
New Feature	Lambda@Edge now enables you to further customize the delivery of content stored in an Amazon S3 bucket, by allowing you to access additional whitelisted headers, including custom headers, within origin-facing events. For more information, see these examples showing personalization of content based on viewer location and viewer device type .	March 20, 2018
New Feature	You can now use Amazon CloudFront to negotiate HTTPS connections to origins using Elliptic Curve Digital Signature Algorithm (ECDSA). ECDSA uses smaller keys that are faster, yet, just as secure, as the older RSA algorithm. For more information, see Supported SSL/TLS Protocols and Ciphers for Communication Between CloudFront and Your Origin and About RSA and ECDSA Ciphers .	March 15, 2018
New Feature	Lambda@Edge enables you to customize error responses from your origin, by allowing you to execute Lambda functions in response to HTTP errors that Amazon CloudFront receives from your origin. For more information, see these examples showing redirects to another location and response generation with 200 status code (OK) .	December 21, 2017
New Feature	A new CloudFront capability, field-level encryption, helps you to further enhance the security of sensitive data, like credit card numbers or personally identifiable information (PII) like social security numbers. For more information, see Using Field-Level Encryption to Help Protect Sensitive Data (p. 219) .	December 14, 2017
Doc history archived	Older doc history was archived.	December, 2017

AWS glossary

For the latest AWS terminology, see the [AWS glossary](#) in the *AWS General Reference*.