

Edgecast

Application Delivery Network Administration Guide

edgecast

Disclaimer

Care was taken in the creation of this guide. However, Edgecast cannot accept any responsibility for errors or omissions. There are no warranties, expressed or implied, including the warranty of merchantability or fitness for a particular purpose, accompanying this product.

Trademark Information

EDGECAST is a registered trademark of Edgecast Inc.

About This Guide

Application Delivery Network Administration Guide

Version 2.50

12/6/2021

© 2021 Edgecast Inc. All rights reserved.

Table of Contents

Introduction	1
Overview	1
Choosing a Data Source	2
Caching & Content Delivery	2
CDN URL	4
Edge CNAME URL	5
How Requests are Handled	6
Phase 1: Request	7
Phase 2: Security	7
Phase 3: Check for Cached Content (POP)	9
Phase 4: Fetch and Deliver	10
Phase 5: Cache	11
Common HTTP Status Codes	12
Configuring an Origin Server	13
Overview	13
Customer Origin Server	14
Origin Configuration	15
Directory Name	15
Domain – IP Address Configuration	16
HTTP Host Header	18
Load Balancing	18
ADN Gateway	19
Customer Origin Optimization	21
Customer Origin Management	21
Firewall Access	24
Best Practices (Dynamic Application)	25

CDN Origin Servers.....	25
Masking the URL of a CDN Origin Server (CNAME).....	26
Setting up DNS for an Edge CNAME Deactivation	28
Edge CNAME Deactivation	28
Securing Your Content	29
Overview	29
Secure HTTP Requests (HTTPS).....	29
Certificate Provisioning System	29
Legacy HTTPS	36
Token-Based Authentication	38
Country Filtering	38
Administering Country Filtering Configurations	39
Protecting Assets Using Both Country Filtering and Token-Based Authentication	42
Cache Management.....	43
Default Cache Management	43
Loading Assets.....	44
Automatically Loading Assets through the Web Services REST API	46
Purging Assets	47
Automated Purging Through the Web Services REST API	47
Manually Purging Assets	48
Purge and Load History	53
Query String Caching.....	54
Query String Logging.....	55
Asset Compression.....	55
Customer Origin Server Compression.....	55
Edge Server Compression	59
Overwriting Cache Headers	61
Glossary.....	63

Introduction

Overview

Web applications deal with information generated from a database or they are driven by user interaction. The dynamic nature of this type of content makes caching impractical, since an asset previously generated by a web application would never be requested again. A major tool in the arsenal of a CDN is to cache assets close to the clients that are requesting it and thereby bypassing the data retrieval from an origin server for subsequent requests. Since caching is impractical for web applications, they cannot take full advantage of the capabilities provided by our CDN. It is in this type of environment where Application Delivery Network (ADN) excels.

Although dynamic content cannot be effectively cached, ADN finds other ways through which to speed up data delivery to your clients. It is able to drastically reduce the amount of time that it takes to communicate with the origin server by optimizing each of the following items:

- Network path
- Server communication
- RFC-compliant protocol communication

The above optimizations allow our dedicated servers to transmit dynamic data at a much faster rate than traditional Internet data routes or even our CDN service can accomplish.

In addition to providing faster web page loads for your existing customers, ADN allows you to expand to far-off markets that were previously out of reach due to website abandonment. These customers that would previously experience prohibitively long web page load time can now load your web pages at blazing speeds.

If your web application also contains static content, then you have the choice to either accelerate that content through ADN or to take advantage of our CDN services (i.e., HTTP Large or the HTTP Small platforms). Using ADN for all of your content will ensure that your entire site is optimally configured for data delivery. Alternatively, using our CDN services to deliver static content allows you to take advantage of caching to reduce the distance that your data must travel without incurring the additional cost of premium ADN service. Either way, our services create an environment where you can accelerate your entire site and thus produce faster load times for your clients.

Choosing a Data Source

Before assets can be served at an accelerated rate to your customers, they must first be retrieved from a web server. You can either retrieve assets from your own web server (i.e., customer origin server) or on a storage server in our CDN (i.e., CDN origin server). Below you will find a brief discussion on the advantages for each type of origin server.

A typical ADN configuration involves a customer origin server. This is due to the fact that your dynamic content is typically generated by code processed by your own web server. Since your content is constantly being updated, it makes sense to point the ADN Gateway directly to your customer origin server.

Content can also be retrieved from a CDN origin server. The main advantage of a CDN origin server is that this type of server is inside our CDN network. Your data is no longer susceptible to network issues due to denial of service attacks or high traffic volume. This guarantees that the data will be available for distribution to your clients. Additionally, since the data does not have to be retrieved from a location outside our network, it can be delivered even faster to your clients. However, this type of configuration is only recommended for static content. Taking advantage of ADN to accelerate static content will ensure that your entire site is quickly delivered to your clients.

Caching & Content Delivery

The URL used to access your content determines how your content will be delivered. Therefore it is important to always request dynamic content using a CDN or edge CNAME URL that is specific to the ADN platform. These types of URLs are listed on the **Application Delivery Network** page. Load this page by finding the **ADN** menu and then selecting **Application Delivery Network**. Using any other type of URL will prevent those referenced assets from being accelerated through ADN.

Tip: You can choose to mix-and-match the delivery system (i.e., ADN, CDN, or traditional Internet data routes) that will be used to serve your web application to your clients. If you choose to serve content from a customer origin server using both our ADN and CDN services, then you will need to create a customer origin on each platform (i.e., ADN, HTTP Large, and/or HTTP Small) through which data will be delivered.

Due to the dynamic nature of your content, it will not be cached by default on an ADN Gateway or an edge server. However, the instructions found in each asset's Cache-Control and Expires headers determine how the asset will be cached by a user agent (e.g., browser). A web browser's cache allows a user to revisit a page without having to wait for it to be transferred from a CDN or customer origin server. This may not be desired for dynamic content, since it could lead to situations where unique content is being cached and/or old content is being served. You should verify that your web server assigns a caching policy that makes sense for

dynamic content. For more information on how assets are cached, please refer to the **Cache Management** chapter.

As mentioned above, there are two types of URLs that can be used to request content through our CDN, which are:

- **CDN URL:** A system-defined URL that identifies a platform/service and a location/asset that will be routed through our network.
- **Edge CNAME URL:** A URL that uses an edge CNAME and a CNAME record to mask a CDN URL. This type of URL is more user-friendly than a CDN URL.

Both types of URLs can be used to identify an asset that will be routed through our content delivery network. The main difference between these two URLs is that an edge CNAME URL allows you to use a custom domain (e.g., www.mydomain.com) to reference content that will be routed through our network.

Important: CDN and edge CNAME URLs are case-sensitive. Using an improper case in the URL will result in a 404 – Not Found error message.

Tip: View a base URL for both CDN and edge CNAME URLs for the ADN platform from the **Application Delivery Network** page. Load this page by finding the **ADN** menu and then selecting **Application Delivery Network**. If you have purchased the SSL Traffic feature for the ADN platform, then HTTPS URLs will also be available from this page.

Note: Our ADN technology leverages our existing CDN to accelerate your website. As a result, URLs that take advantage of ADN are also referred to as CDN and edge CNAME URLs.

In the following example, an edge CNAME called "www.mydomain.com" has been set to mask the following segment of a CDN URL "adn.0001.edgecastcdn.net/800001/Business." Both of the following URLs request the same asset (i.e., Article.php).

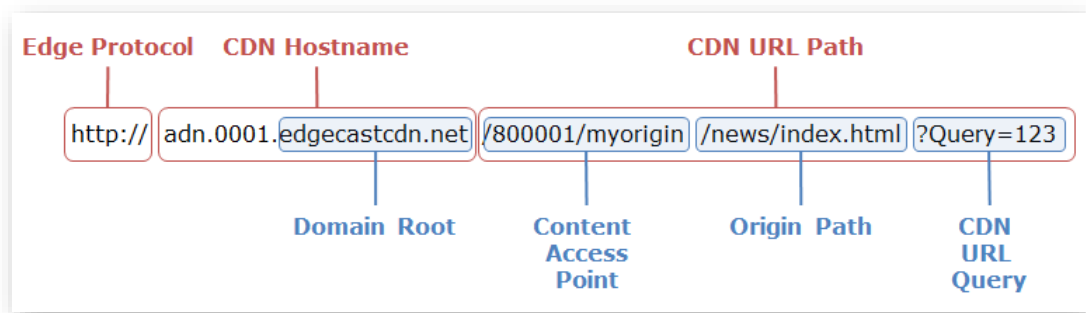
URL Type	URL
CDN URL	http://adn.0001.edgecastcdn.net/800001/Business/Documents/Article.php
Edge CNAME URL	http://www.mydomain.com/Documents/Article.php

CDN URL

A CDN URL is any URL that points to our CDN and contains a content access point. Each customer account contains at least one CDN URL per platform. A CDN URL will be assigned to your account for the following items:

- CDN storage (root folder)
- Each customer origin configuration (root folder)
- Each edge CNAME configuration

The following diagram indicates the different components in a CDN URL. Keep in mind that xxxx represents your CDN account number.

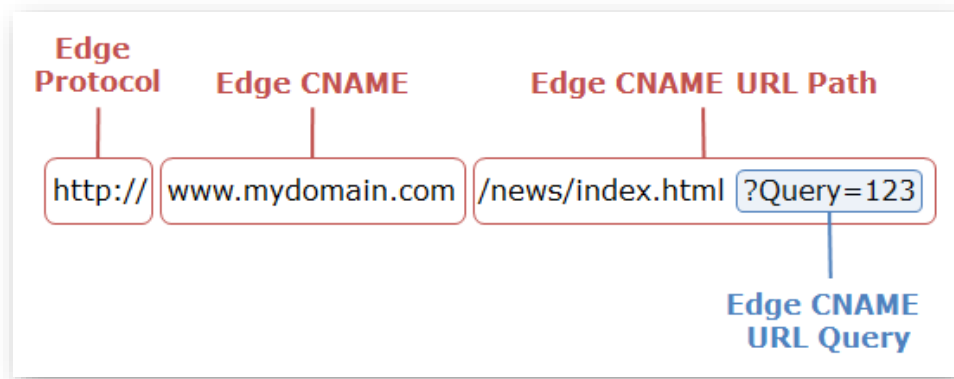


Name	Description
Edge Protocol	The edge protocol is always either "HTTP" or "HTTPS" for the ADN platform.
CDN Domain	The CDN domain is the unique domain used by our CDN to determine the platform and the CDN customer associated with a request. In the above example, the CDN domain is "adn.xxxx.edgecastcdn.net," where "adn" identifies the ADN platform and xxxx stands for your CDN account number.
CDN URL Path	The CDN URL path is the portion of the URL that appears after the CDN domain. The CDN URL query portion of the URL is optional.
Domain Root	The domain root identifies the base CDN domain that has been assigned to your CDN account. In the above example, the domain root is edgecastcdn.net.

Name	Description
Content Access Point	<p>The content access point, which appears directly after the CDN domain, allows our CDN to identify the source for the requested content. A content access point can be one of the following for the ADN platform:</p> <ul style="list-style-type: none"> • CDN origin server: /00xxxx • Customer origin server: /80xxxx/<i>CustomerOrigin</i> <p>The two digit number that precedes the account number (i.e., xxxx) is called the origin identifier. For the purposes of the ADN platform, it instructs our CDN to either route traffic from a CDN or customer origin server. The term <i>CustomerOrigin</i> represents the name assigned to the desired customer origin configuration.</p>
Origin Path	<p>The origin path indicates the location of the desired asset. This path starts from the root folder of the origin server identified by the content access point.</p>
CDN URL Query	<p>Optional. A query string can be appended at the end of the CDN URL. By default, query strings are ignored for the purposes of caching and event logging.</p>

Edge CNAME URL

An edge CNAME URL is any URL that uses an edge CNAME to mask a CDN URL. The following diagram indicates the different components in an edge CNAME URL.



Name	Description
Edge Protocol	The edge protocol is always either "HTTP" or "HTTPS" for the ADN platform.
Edge CNAME	There is a CDN and a DNS aspect to an edge CNAME. The CDN aspect is an edge CNAME configuration that allows our servers to identify a contact access point. In addition to this configuration, a DNS server must be configured to recognize the alias for the CDN URL. The alias set on the DNS server should match the domain (e.g., adn.xxxx.edgecastcdn.net) associated with the edge CNAME configuration. If a directory path has been specified in the edge CNAME configuration, then our servers will rewrite the URL to point to the appropriate folder.
Edge CNAME URL Path	The edge CNAME URL path indicates the location of the desired asset on the origin server identified by the edge CNAME. The starting point for this path is determined by the relative path associated with the corresponding edge CNAME configuration. If a relative path has not been defined in the edge CNAME configuration, then the edge CNAME URL path starts from the root folder of the origin server.
Edge CNAME URL Query	Optional. A query string can be appended at the end of the edge CNAME URL. By default, query strings are ignored for the purposes of caching and event logging.

How Requests are Handled

Whenever an asset is requested for the ADN platform, it will undergo the following procedure:

1. Request
2. Security
3. Check for Cached Content (POP)
 - Check for Content Freshness (POP)
4. Fetch and Deliver
5. Cache

Below you will find an explanation for each phase in this procedure.

Phase 1: Request

Before an asset can be delivered to a client, a request must be made to an edge server that resides on one of our POPs. Making a request to one of our edge servers is simply a matter of using a CDN or edge CNAME URL that points to the desired asset. The following sample CDN URL points to an asset that will be delivered through the ADN platform:

```
http://adn.0001.edgecastcdn.com/800001/Server/index.html.
```

This request will be handled by the point-of-presence (POP) closest to that client. This ensures that a minimal amount of time will be required to establish a connection between the client and the edge server.

Phase 2: Security

The first check performed by an edge server in a POP is to find out whether the requested content has been secured. We offer the following methods for securing your content:

- Country Filtering
- Token-Based Authentication
- Rate Limiting
- HTTP Rules Engine
- Web Application Firewall

If the request fails to meet any of the above security requirements, then our edge servers will deny the request. No additional steps will be taken.

Country Filtering

Country Filtering allows you to secure content by directory. If the requested content resides in a secured directory, then the location of the user requesting the content determines whether the request will be authorized. If the request is denied, then the HTTP client will receive a 403 Forbidden status code.

Reminder: Country Filtering is applied recursively to a secured location. For example, an asset that is located in a subfolder of a secured directory will still be protected by Country Filtering.

Rate Limiting

Rate Limiting restricts the number of requests that the CDN may forward to your web servers. Once the specified rate is exceeded, requests will be throttled to the specified rate. Requests that exceed this rate may either be dropped, redirected, sent a custom response, or simply be logged.

Token-Based Authentication

If the requested content has been secured through Token-Based Authentication, then an edge server will look for a token value in the query string. If it finds one, then the token value will be decrypted according to the encryption key(s) specified for the ADN platform. A check will then be performed to see whether the HTTP client (e.g., browser) meets the specified requirements.

Note: Content can be secured by folder location or through HTTP Rules Engine. For more information, please refer to either the **Token-Based Authentication Administration Guide** or the **HTTP Rules Engine Administration Guide**.

By default, an HTTP client will receive a 403 Forbidden status code when any of the following is true for an asset protected by Token-Based Authentication:

- The CDN or edge CNAME URL did not include a token value in the query string.
- The specified token was invalid. This occurs for tokens generated using an invalid encryption key or an invalid parameter was used when generating an encryption key.
- The HTTP client did not meet the token requirements.

Tip: Instead of returning a 403 Forbidden status code for the cases mentioned above, you can take advantage of custom denial handling to redirect users or to report a different HTTP status code.

HTTP Rules Engine

HTTP Rules Engine is a powerful tool that can be used for a variety of purposes. One feature provided in this tool is "Deny Access (403)." This feature will return a 403 Forbidden status code to the HTTP client whenever a user-defined set of conditions are met. For information on how you can define the conditions (i.e., match options) that will be used to deny access to sensitive content, please refer to the **HTTP Rules Engine** guide.

Web Application Firewall

Web Application Firewall (WAF) is designed to secure an origin server against malicious attacks. Once WAF has been configured, HTTP Rules Engine determines the types of requests that will be secured by WAF. If a request is screened by WAF, then a WAF profile determines whether it will be identified as a threat. A WAF instance determines whether detected threats will be allowed or blocked. WAF returns a 403 Forbidden for all blocked requests.

Phase 3: Check for Cached Content (POP)

If the requested asset passes all of the above security requirements, then a check is performed to see whether the asset is currently cached on the POP handling the request.

Reminder: Due to the dynamic nature of content served from the ADN platform, the caching mechanism is bypassed by default. As a result, a cached version of your content should not be found.

Check for Content Freshness (POP)

If a cached version of the requested asset is found on the POP handling the request, then a check for content freshness will be performed.

Note: If the cached version of the requested asset contains a "Cache-Control: must-revalidate" header value, then our edge and ADN gateway servers will honor it. This means that they will always check the origin server for a newer version of a stale asset with this header value.

Check for content freshness:

1. An asset's time to live (TTL) will be calculated through its Cache-Control and Expires headers. If the asset is still fresh, then it will be immediately delivered to the client and no further steps will be performed.
2. If the asset has exceeded its TTL, then the POP will check the ADN Gateway server for cached content.
 - If a cached version of the asset is found, then it will be checked for freshness.
 - If the TTL is valid, then the following will take place:
 - The asset will be served to the client via an edge server corresponding to the POP closest to the client.
 - It will then be cached on that edge server.
 - If the TTL is expired, then the cached version on the ADN Gateway server will be compared against the customer origin server.
 - If the same version of the asset is found on the customer origin server, then the following will take place:
 - a. The headers on the ADN Gateway server will be updated.
 - b. The asset will be served to the client via an edge server corresponding to the POP closest to the client.
 - c. It will then be cached on that edge server.

- If a newer version is found, then the following will take place:
 - a. The ADN Gateway server will receive a new version of the asset.
 - b. The ADN Gateway server will forward it to an edge server corresponding to the POP closest to the client.
 - c. It will then be cached on that ADN Gateway server.
 - d. The edge server will deliver it to the client.
 - e. It will then be cached on that edge server.
- If cached content is not found, then the following will take place:
 - i. The asset will be requested from the customer origin server and an ADN Gateway server will receive it.
 - ii. The ADN Gateway server will forward it to an edge server corresponding to the POP closest to the client.
 - iii. It will then be cached on that ADN Gateway server.
 - iv. The edge server will deliver it to the client.
 - v. It will then be cached on that edge server.

Phase 4: Fetch and Deliver

A new version of the requested content will need to be retrieved from an ADN gateway server under either of the following conditions:

- The requested content is not cached on the edge server handling the request.
- A cached version of the requested content was found but its TTL has expired.

Either of the above conditions will cause an edge server on the respective POP to request the asset from an ADN gateway server. At which point, a new version of the requested content will need to be retrieved from the origin server under either of the following conditions:

- The requested content is not cached on the ADN gateway server handling the request.
- A cached version of the requested content was found but its TTL has expired.

Either of the above conditions will cause an ADN gateway server to request the asset from a CDN or customer origin server. Once it starts receiving the asset from the origin server, an ADN gateway server will immediately start delivering it to the client that requested it via the edge server that forwarded the initial request.

Reminder: A default ADN configuration does not cache content and therefore a cached version of the requested content should not be found.

Phase 5: Cache

Once the entire asset has been retrieved from the origin server, it will be cached on the corresponding ADN gateway server and edge server provided that all of the following conditions are met:

Condition	Description
1 st or 2 nd Hit Caching	By default, content must be requested twice before it is cached on an ADN gateway server or an edge server.
Cache Policy	The Cache-Control and Expires headers provided in the origin server's response determine whether the requested content will be cached. If those response headers are missing, then it will be cached with a max-age of 7 days.
Bypass Cache	Rule Engine's Bypass Cache feature must be disabled for the requested content. For more information on the Bypass Cache feature, please refer to the Rules Engine Administration Guide.

Note: Rules Engine, which must be purchased separately, can be used to customize how our CDN caches assets. For example, our CDN may be configured to cache an asset even if a user aborts his/her download.

Key information:

- Ensure that dynamic content is not cached by a user agent (e.g., browser) by configuring your web server to add a "Pragma: no-cache" or a "Cache-control: no-cache" directive to the response generated for that content.
- Override an asset's cache headers and prevent it from being cached by specifying "no-cache" as a query string parameter when requesting it.
 - i. **Example:** An asset requested using the following URL will not be cached:
`http://adn.0001.edgecastcdn.net/800001/Images/Logo.png?no-cache`

Common HTTP Status Codes

A brief explanation is provided below for the common HTTP status codes that may be returned by a CDN or customer origin server.

HTTP Status Code	Description
200 OK	An origin server returns this status code when it authorizes a request for an asset that is not currently cached on the POP serving the region from which the request originated or the ADN Gateway server.
302 Found	An origin server returns this status code when a URL redirection for the requested asset was detected. This status will also be returned for cookie revalidations.
304 Not Modified	An origin server returns this status code when it authorizes a request for a stale asset that has not been modified since it was cached. This status code will be returned regardless of whether the asset was cached on the POP serving the region from which the request originated or an ADN Gateway server. The asset will be served from either the ADN Gateway server or the edge server that handled the request. The Cache-Control and Expires headers for that asset will be updated.
403 Forbidden	<p>An origin server or an edge server returns this status code when an unauthorized request is received. The requestor will be denied access to the requested asset. Additionally, a check will not be performed on the origin server to verify the freshness for the cached version of the requested asset.</p> <p>An edge server will return this status code when the request for an asset protected by Token-Based Authentication does not include a valid token or if the requestor does not meet the conditions specified in the included token.</p>
404 Not Found	An origin server returns this status code when the requested asset was not found.
504 Gateway Timeout	An edge server returns this status code when it is unable to communicate with an origin server when trying to revalidate the requested asset.

Configuring an Origin Server

Overview

Before assets can be served through our CDN to your customers, they must first be retrieved from a web server. This type of server is known as an origin server, since it is the origin or source of the content that will be delivered to your users. A brief description is provided below for each type of origin server.

Origin Server Type	Description
Customer Origin Server	<p>Our CDN service can be configured to recognize a web server that is external to our network (e.g., web hosting). The two main advantages of using a customer origin server are:</p> <ul style="list-style-type: none">• Minimal Setup: Setting up a customer origin server is as simple as creating a customer origin configuration from the MCC. All of your data can remain as it is currently structured on your web server. This can save you the amount of time it takes to upload and maintain that data on a separate server.• Response Headers: Our edge servers will forward the majority of response headers provided by your web server to your users. A benefit from this behavior is that you can define the cache instructions that will be assigned to your assets on your web server without having to purchase HTTP Rules Engine.
CDN Origin Server	<p>Our CDN storage service allows you to upload content to a CDN web server dedicated to file storage. This type of server is known as a CDN origin server. Content in CDN storage can then be distributed to your users via our network. Our CDN origin servers are located within our CDN network which provides the following main benefits:</p> <ul style="list-style-type: none">• Protection: Our highly distributed network provides protection against network issues due to denial of service attacks or high traffic volume. This guarantees that the data will be available for distribution to your clients.• Speed: Data does not have to be retrieved from a location outside our network. This allows it to fully leverage our high-speed network to deliver content faster to your clients.

Customer Origin Server

Content stored or generated from a third-party web server (e.g., web hosting) can be delivered through our network. This requires that you define a customer origin configuration. This type of configuration maps one or more servers to a CDN URL.

Key information:

- The maximum number of customer origin configurations per platform is 100.
 - Each configuration is specific to a platform.
 - A directory name must be specified when defining a customer origin configuration. This directory name identifies your customer origin configuration in a CDN URL.
 - Since CDN URLs are case-sensitive, make a note of the case used when defining the directory name.
 - An edge CNAME configuration and a CNAME record may be used to generate a user-friendly URL known as an edge CNAME URL. For more information, please refer to the **Masking the URL of a CDN Origin Server (CNAME)** section.
 - Define the set of servers that will fulfill HTTP requests.
 - If the SSL Traffic feature has been enabled on the platform in question, then you may also define a set of servers that will fulfill HTTPS requests.
2. A customer origin configuration consists of the following settings:
- Origin Configuration
 - Directory Name
 - Domains/IP Addresses
 - HTTP Host Header
 - Load Balancing
 - ADN Gateway

Origin Configuration

Domains associated with a customer origin configuration must be resolved to an IP address before a request can be served to it. The **Origin Configuration** option determines whether our servers will prefer to resolve a domain to an IPv4 or IPv6 address.

Setting	Description
Default	Indicates that domains should be resolved to IPv4 addresses only. <hr/> Note: We reserve the right to change the behavior of this default setting at any time. <hr/>
V6 Preferred Over V4	Indicates that our edge servers will resolve domains to IPv6 addresses whenever possible. If an IPv6 address for that domain does not exist, then the domain will be resolved to an IPv4 address.
V4 Preferred Over V6	Indicates that our edge servers will resolve domains to IPv4 addresses whenever possible. If an IPv4 address for that domain does not exist, then the domain will be resolved to an IPv6 address.
V4 Only	Indicates that domains should be resolved to IPv4 addresses only.
V6 Only	Indicates that domains should be resolved to IPv6 addresses only.

Directory Name

The **Directory Name** option defines how your customer origin configuration will be identified in a CDN URL. This name is specified as a part of the content access point (i.e., /80xxxx/**directory**). Our edge servers interpret the content access point to identify the starting location for your assets.

For example, if the primary purpose of your origin server is to serve images, then you might create a customer origin configuration whose **Directory Name** option is set to "images." A sample CDN URL that may be used to access assets on this customer origin server is provided below.

```
http://adn.xxxx.edgecastcdn.net/80xxxx/images
```

Note: The term "xxxx" represents your CDN account number.

The above sample CDN URL points to the root folder on the server(s) associated with the "images" customer origin configuration. Append the desired relative path to the content that you would like to request. For example, the following sample CDN URL points to this location "http://www.myserver.com/photography/clientX" on your customer origin server:

```
http://adn.xxxx.edgecastcdn.net/80xxxx/images/photography/clientX
```

Tip: If you would like to use a more user-friendly URL, then you should take advantage of an edge CNAME.

Note: CDN and edge CNAME URLs are case-sensitive.

Domain – IP Address Configuration

A customer origin configuration must point to one or more web servers. These web servers are defined through the **Hostname Configuration** section. The **HTTP Edge Protocol** and the **HTTPS Edge Protocol** options in that section define the set of servers that can fulfill HTTP or HTTPS requests, respectively.

Note: The ability to define a set of web servers that can fulfill HTTPS requests requires the SSL Traffic feature for the ADN platform. This feature, in combination with an SSL certificate and an edge CNAME, allows the use of an HTTP Secure (HTTPS) URL when delivering content over our network. Please contact your CDN account manager to activate this feature.

Key information:

- Use one of the following formats to define a web server:

Format	Example
<i>protocol://domain:port</i>	http://www.mydomain.com:80
<i>protocol://IPv4Address:port</i>	http://10.10.10.255:80
<i>protocol://[IPv6Address]:port</i>	http://[1:2:3:4:5:6:7:8]:80

- An optimal configuration requires that the servers specified for each customer origin reside in relatively close vicinity to one another. If you would like to specify one or more servers that are located in a different geographic region, then we highly recommend that you create a separate customer origin configuration for those servers.
- In order to avoid DNS latency when loading content from a customer origin server, our edge servers proactively resolve each domain defined in a customer origin configuration.
- Our edge servers respect the DNS TTL defined for an origin server when resolving a domain.
- If our edge servers are unable to resolve a domain or if the domain's DNS TTL has expired, then they will reattempt to resolve it at regular intervals until it is resolved. Since this action may incur additional charges by your DNS provider, it is recommended that all domains associated with a customer origin configuration be registered in DNS with a reasonable TTL (e.g., 3600 seconds or more).
- Any combination of domains and IP addresses may be used to identify the set of origin servers that can honor requests for the customer origin being configured.

- The set of domains and IP addresses defined for either the **HTTP Edge Protocol** or the **HTTPS Edge Protocol** list must be assigned the same protocol (i.e., HTTP or HTTPS).
- Defining more than one server for either the **HTTP Edge Protocol** or the **HTTPS Edge Protocol** list will cause requests to be load balanced across those origin servers for that protocol.
- A maximum of 10 domains and/or IP addresses per protocol (i.e., HTTP or HTTPS) can be associated with a customer origin configuration.

Reminder: A domain must be resolved to an IP address before a request can be forwarded to it. The **Origin Configuration** option defines how domains will be resolved to an IP address.

HTTP Requests

The **HTTP Edge Protocol** option defines the set of web servers that can handle HTTP requests.

Key information:

- Make sure that all web servers defined for the **HTTP Edge Protocol** option use the same protocol (i.e., either HTTP or HTTPS).
- If a valid origin server has not been specified under the **HTTP Edge Protocol** option, a 404 – Not Found error message will be returned to the user when content is requested from your origin server using the HTTP protocol.

HTTPS Requests

The SSL Traffic feature enables an additional customer origin configuration option called "HTTPS Edge Protocol." This option functions in the same way as the **HTTP Edge Protocol** option, except that it determines how HTTPS requests are handled.

Key information:

- The HTTP protocol may be used to fulfill HTTPS requests, but this type of configuration would sacrifice end-to-end encryption of the request and response. If this describes your preferred configuration, please make sure that all servers are identified using the HTTP protocol.
- If a valid domain has not been specified under the **HTTPS Edge Protocol** option, a 404 – Not Found error message will be returned to the user when content is requested from your origin server using the HTTPS protocol.
- Base CDN and edge CNAME URLs can be viewed from the **SSL URL** section on the ADN home page.

HTTP Host Header

HTTP 1.1 requires a Host header to be sent with each request. A Host header identifies the domain/IP address and port associated with a request. This is especially useful when there are multiple virtual domains hosted on a single physical server or load-balanced set of servers.

Each customer origin configuration allows a Host header value to be configured. It is recommended that you set this value to either one of the domains/IP addresses from the HTTP/HTTPS hostname lists, or else to an edge CNAME that references this customer origin server.

Note: By default, the **HTTP Host Header** option is set to the first domain or IP address defined under the **Hostname Configuration** section.

Note: If the **HTTP Host Header** option is set to blank, then the request URI determines the value for the Host request header.

Load Balancing

The domains/IP addresses associated with the **HTTP Edge Protocol** option are resolved into a list of IP addresses through DNS. A separate list is generated for the **HTTPS Edge Protocol** option. If multiple unique IP addresses are associated with either option, then the selected load balancing mode determines which customer origin server will handle the next request.

The available load balancing options are:

- **Round Robin:** This mode will balance requests between all of the servers listed for a particular protocol. In other words, it will send the first request to the first server on the list. The next request will be sent to the next server on the list.
- **Primary & Failover:** This mode indicates that the specified servers form an ordered failover list. In other words, all requests will be forwarded to the first server on the list. If that server is unavailable, then the request will be sent to the next server on the list. This process continues until a server is able to honor the request. A server is unavailable when a TCP connection is refused or if the connection times out. The order of the failover sequence can be controlled by rearranging entries through the up/down double arrow buttons.

Note: These load-balancing options are completely independent from any load balancing that may already exist on the customer origin server. For instance, traffic for a single IP address might be balanced across several physical servers.

ADN Gateway

ADN Gateway servers are responsible for ensuring efficient data delivery between the CDN and your customer origin servers. Before an ADN Gateway server can assume this responsibility, our network will need to calculate the top three ADN Gateway servers that can provide the best performance for a customer origin configuration. This is calculated by retrieving an asset from each web server associated with a customer origin configuration.

Key configuration information:

- **File Size:** The ideal type of asset to test the latency between ADN Gateway servers and customer origin server(s) is one whose size is approximately 5 – 10 KB.
- **Sample Asset:** A sample performance test asset can be found at:
 - https://my.edgecast.com/images/PerformanceTestObject_5k.gif
- **Upload:** Upload the sample asset or one of your own choosing to each server specified in your customer origin configuration. If you are uploading an asset to multiple servers, make sure to always use the same relative path from the root folder.
- Set the **Validation Path** option to a URL that points to the desired asset.
 - Make sure that this URL uses the HTTP host header associated with your customer origin configuration.
- Click **Validate** to check the communication with each specified server and to verify that the asset is available on each of those servers.
 - **Error:** If an error is reported, make sure to address the error and then revalidate the asset.
 - **200 OK:** If the test returns 200 OK for all specified domains/IP addresses, click **Add** or **Update**. This action will create or update the customer origin configuration and cause all ADN Gateway servers to retrieve the specified asset from each server specified within the customer origin configuration. The top three ADN Gateway servers that provide the best performance will then be assigned to the customer origin.

The following sample scenario provides an example as to how the **Validation Path** option should be configured.

Item	Scenario Description
HTTP Host Header	A customer origin server's HTTP host header is set to: <ul style="list-style-type: none">• adn.mydomain.com
Sample Asset Configuration	A sample asset has been uploaded to all of the servers defined under the Hostname Configuration section. The relative path to this sample asset is: <ul style="list-style-type: none">• /Permanent/Background.png
Validation Path	The Validation Path option should be set to: <ul style="list-style-type: none">• http://adn.mydomain.com/permanent/background.png

Important: If a sub-optimal configuration is detected, then a message will appear directly below the **ADN Gateways** section. It is highly recommended that you verify that the detected issue will not affect the acceleration of your website. For example, if ADN detects that your site is redirecting URLs to another server, then this could lead to a situation where the wrong ADN Gateway server is being used to accelerate your website.

Note: If the asset defined in the **Validation Path** option needs to be moved, renamed, or deleted, make sure to perform the same action on all servers specified in your customer origin configuration and then update the **Validation Path** option to reflect this change.

Note: Use the **Follow redirects** option to determine whether our edge servers will respect a URL redirect when validating the set of optimal ADN gateway servers for a customer origin configuration.

Note: Creating a customer origin will automatically revalidate the optimal path between the customer origin and ADN Gateway servers.

Note: The use of the **Reselect** option is recommended when a customer origin has been modified or after an ADN Gateway error has been resolved.

Note: A new ADN Gateway configuration may take up to an hour to take effect. During this time period, the ADN Gateway servers previously assigned to a customer origin configuration will continue to accelerate your site.

Customer Origin Optimization

In addition to creating a customer origin configuration for the ADN platform, it is recommended that you ensure that your customer origin server has been optimized for use with ADN. This primarily involves configuring your customer origin server to keep connections alive. There are two ways through which a customer origin server can maintain a connection alive, which are through a web server and through the OS. Since the default configuration for each type of web server and OS varies, it is important that you verify that this feature is turned on.

Tip: You should tailor the number of requests that each connection should handle according to the amount of traffic that normally flows through your web application during the specified time period for which your connection will stay alive (e.g., 30 minutes).

Note: If you have disabled TCP and HTTP persistent connections due to security reasons, then an alternative would be to create a virtual web server on which this feature has been enabled. This virtual web server should only have a single open port and it should only accept connections from the IP ranges listed on the **Customer Origin** page. Load this page by finding the **ADN** menu and then selecting **Customer Origin**. Make sure that your customer origin points to the virtual web server.

Customer Origin Management

A customer origin configuration is required to serve content from your servers to your users via our CDN. This section explains how to create, modify, and delete customer origin configurations.

Creating a Customer Origin Configuration

This section provides step-by-step instructions on how to create a customer origin configuration.

Key information:

- A customer origin configuration must point to a DNS record that has been fully propagated by your DNS provider. This may be checked through the use of the "dig" command-line tool, which is a DNS query tool that allows you to query your DNS provider for the DNS record that corresponds to your domain. When performing this dig, it is recommended that you use the "+trace" parameter. This will ensure that the DNS provider provides a direct response on the DNS record for the specified domain. The syntax for this command is provided below.

```
dig domain +trace
```

In the bottom section of the response, you should see an A record that points your domain to an IP address. This indicates that your DNS record has been fully propagated. An excerpt from a sample response is shown below.

```
www.mycustomerorigin.com.    3600    IN      A       10.10.10.101
mycustomerorigin.com.        3600    IN      NS      dns02.dnsauthority.com.
mycustomerorigin.com.        3600    IN      NS      dns01.dnsauthority.com.
;; Received 100 bytes from 100.100.100.101#53(100.100.100.101) in 20 ms
```


- As a precautionary measure, it is recommended that you set the DNS TTL for your domain to a low value until you have confirmed that our edge servers can properly communicate with your origin server. A low DNS TTL reduces the amount of time that an improper DNS configuration will affect your origin server, while increasing the number of DNS queries sent to your DNS provider.
- The required options for a new customer origin are a folder name, a domain/IP address, and an HTTP host header. For your convenience, the HTTP host header is automatically populated when you add the first domain/IP address.
- Make sure that your origin server does not restrict access to the IP address blocks listed on the **Customer Origin** page.
- It may take up to an hour for your new customer origin configuration to take effect.

To create a customer origin configuration

1. Navigate to the **Customer Origin** page by finding the **ADN** menu and then selecting **Customer Origin**.
2. In the **Directory Name** option, type the name of the folder that will be associated with the desired customer origin server. This folder will become a part of the contact access point in the CDN URL (e.g., `http://adn.0001.edgecastcdn.net/800001/FolderName`).
3. Perform one of the following:
 - If you would like customers to access the content of the customer origin server being configured through the HTTP protocol, then you should:
 - i. Make sure that the **HTTP Edge Protocol** option is marked.
 - ii. In the **Hostname or IP Address** option, you should type the domain or IP address of the customer origin server followed by a colon and the port through which communication will take place.
 - iii. Click **Add**.
 - iv. Repeat this step until you have finished adding all of the domains that will be associated with this customer origin for `http://` requests.

- If you would like customers to access the content of the customer origin server being configured only through the HTTPS protocol, then the **HTTP Edge Protocol** option should be cleared and you should proceed to the next step.
4. If you have purchased the SSL traffic feature for this platform, perform one of the following:
 - If you would like customers to access the content of the customer origin server being configured through the HTTPS protocol, then you should:
 - i. Make sure that the **HTTPS Edge Protocol** option is marked.
 - ii. In the **Hostname or IP Address** option, you should type the domain or IP address of the customer origin server followed by a colon and the port through which communication will take place.
 - iii. Click **Add**.
 - iv. Repeat this step until you have finished adding all of the domains that will be associated with this customer origin for https:// requests.
 - If you would like customers to access the content of the customer origin server being configured only through the HTTP protocol, then the **HTTPS Edge Protocol** option should be cleared and you should proceed to the next step.
 5. Verify or set the **HTTP Host Header** option to one of the following:
 - A domain specified in either step 3 or 4.
 - An edge CNAME that points to a domain defined in either step 3 or 4.
 - Blank. The request will determine the value assigned to the Host request header.
 6. Upload an asset to each server that has been specified for this customer origin configuration. Our servers use this asset to choose the optimal ADN Gateway servers and data routes for your customer origin configuration. Keep in mind that 5 KB is the optimal file size.
 7. In the **Validation Path** option, specify a URL that points to a sample asset on your customer origin server. The domain specified in this URL should match the one defined in the **HTTP Host Header** option unless that option has been set to blank.
 8. Click **Validate**. If the result returns "200 OK" for all domains/IP addresses, then proceed to the next step.
 9. Click **Add** to save your customer origin configuration. The top three ADN Gateway servers that are best suited for serving content from that customer origin server will eventually be displayed next to the **ADN Gateway** option.


Modifying a Customer Origin Configuration

A customer origin configuration can be modified at any time by clicking the  next to the desired customer origin. The configuration associated with that customer origin will appear. Simply make the desired changes and then click **Update** to apply them.

Note: If an edge CNAME points to a customer origin configuration, then you will not be allowed to modify the name of the folder associated with that customer origin configuration. If you would like to change the folder name, you will need to first delete the associated edge CNAME.

Note: It may take up to an hour for changes to your customer origin configuration to take effect.

Deleting a Customer Origin Configuration


A customer origin configuration can be deleted at any time by clicking the  next to the desired customer origin. Once you have confirmed the deletion, it will be removed from the list.

Note: If an edge CNAME points to a customer origin configuration, then you will not be allowed to delete the associated customer origin configuration. If you would like to delete it, you will need to first delete the associated edge CNAME.

Note: It may take up to an hour for the deletion of your customer origin configuration to take effect.

Firewall Access

For the purpose of fulfilling requests, our edge servers require access to all servers associated with a customer origin configuration. Please ensure that your firewall allows access to all of the IP blocks listed in the **Whitelist IP Blocks** section of the **Customer Origin** page.

Tip: Export a list of the IPv4 and IPv6 blocks that should be whitelisted by clicking  from the **Whitelist IP Blocks** header.

The **Whitelist IP Blocks** section contains a superset of IP addresses that includes:

- The IP blocks defined under the **The following CDN IPs can access your origin** section.
- The IP blocks for future POPs.

Note: Once the IP blocks defined under the **Whitelist IP Blocks** section have been whitelisted on your firewall, it is unnecessary to add the IP blocks defined under the **The following CDN IPs can access your origin** section.

Best Practices (Dynamic Application)

If your customer origin hosts a dynamic application, then it is highly recommended that you do not use a user's IP address to maintain a session instance. This type of configuration is unsupported, since the client does not connect directly to the customer origin server. Instead, the client connects to a server on our network, and then that server connects to a customer origin server. If you would like to maintain a session for your dynamic application, we recommend that you use a cookie to identify the session. For example, a cookie could keep track of a unique ID for each client's session.

CDN Origin Servers

Our CDN storage solution allows the storage of content on our storage servers (i.e., CDN origin servers). Content stored in this manner can also be delivered at a quicker rate by taking advantage of the many optimizations that ADN provides. Keep in mind that these types of servers have already been configured to provide optimal data delivery performance when used in conjunction with ADN. As a result, you can simply upload your assets to the CDN origin server without having to perform any additional configuration.

Content stored on our CDN storage solution can then be made available to clients either through:

- A system-defined URL (i.e., CDN URL)
- An edge CNAME URL

Note: A list of CDN and edge CNAME URLs that may be used to access content in our CDN storage solution may be viewed from the **URLs for CDN Origin** section of the **Application Delivery Network** page.

Key information:

- Update and manage content on CDN storage using any of the following solutions:
 - SFTP protocol
 - rsync protocol
 - FTP protocol (insecure)
- Access content uploaded to CDN storage through the following base CDN URL:

```
http://adn.xxx.edgecastcdn.net/00xxx/
```

- Use an edge CNAME to provide a more user-friendly URL to your clients.

Note: For detailed information on how to upload content to CDN storage, please refer to the **CDN Storage - Data Upload** guide.

Masking the URL of a CDN Origin Server (CNAME)


The URL to the location on a CDN origin server where your assets reside can be masked through the use of a Canonical Name (CNAME) record and an edge CNAME. An edge CNAME informs our edge servers of the CNAME record that you would like to use, the domain used by that CNAME, and an optional path to which URLs that take advantage of the edge CNAME can be rewritten.

Important: Informing our edge servers that you would like to use a CNAME record will not update or set that CNAME record on your DNS server. Before you can take advantage of the desired CNAME, you will need to set a CNAME record on your DNS server.

Example: Instead of linking to assets using the following URL:
`http://adn.0001.edgecastcdn.net/000001/Videos/Technical`, you could create a CNAME record that points to that folder (e.g., `technical.videos.com`). You could then link to the same folder using this user-friendly URL: "`http://technical.videos.com`."

Note: You may also take advantage of an edge CNAME to mask the URL of your customer origin server.

When configuring our edge servers to recognize a CNAME, you should keep in mind all of the following:


- Do not specify a protocol (e.g., `http://`, `https://`, `ftp://`, etc.).
- All alphabetical characters in the CNAME should be specified in lower-case letters.
- You will need to set a CNAME record on your DNS server. The alias set on the DNS server should match the domain (e.g., `adn.xxxx.edgecastcdn.net`) associated with the edge CNAME configuration. If a directory path has been specified, then our servers will rewrite the URL to point to the appropriate folder.
- Adding, modifying, or deleting an edge CNAME may take up to an hour to take effect.
- The Analytics product offers a variety of different ways to view CDN usage activity. However, if you would like to organize report data by edge CNAME, then you will need to enable the **Custom Reports** option on each desired edge CNAME. This will cause our servers to keep track of the amount of data transferred and hits for each selected edge CNAME. A  icon will appear next to each edge CNAME for which reporting has been enabled.

Note: A custom report can be generated by finding the **Analytics** menu and then selecting **Custom Reports**. After which, select the desired report from the side navigation bar. For more information, please refer to the **Custom Reports** chapter in the **EC360 Analytics Suite** guide.


To configure our edge servers to recognize your CNAME

1. Navigate to the **Edge CNAMEs** page by finding the **ADN** menu and then selecting **Edge CNAMEs**.
2. In the **New Edge Cname** option, type the name of the desired CNAME record. The CNAME should be specified in lower-case letters and should not include the protocol (i.e., http://).
3. Select whether the specified CNAME will point to a customer origin or CDN origin server.
4. By the **Points to** option, select the root location on the origin server to which the CNAME will be pointed. If you would like to indicate a specific folder, then you should type a forward slash (/) followed by the path to the desired folder.
5. In the **Custom Reports** option, choose whether to enable custom data logging for the edge CNAME being created.
6. Click **Add**.
7. Make sure that a CNAME record that points to the same domain has been registered on your DNS server. This CNAME record must match the name assigned to your edge CNAME.

To modify an edge CNAME

1. Navigate to the **Edge CNAMEs** page by finding the **ADN** menu and then selecting **Edge CNAMEs**.
2. Click the  next to the edge CNAME that you would like to modify.
3. Make the desired changes.
4. Click **Update** to save your changes.
5. If you have modified the **New Edge Cname** option, make sure to update the CNAME record registered on the DNS server.

To delete an edge CNAME

1. Navigate to the **Edge CNAMEs** page by finding the **ADN** menu and then selecting **Edge CNAMEs**.
2. Click the  next to the edge CNAME that you would like to delete.
3. When prompted, confirm the deletion of the selected edge CNAME.
4. Make sure to update or delete the corresponding CNAME record via a DNS service provider. Please refer to the **Setting up DNS for an Edge CNAME Deactivation** section for more information.

Note: HTTP Rules Engine allows you to match all requests that originate from a particular edge CNAME. However, an edge CNAME cannot be deleted if it is referenced by a rule. In order to delete the edge CNAME in question, please make sure to first modify or delete all match options that reference it.

Setting up DNS for an Edge CNAME Deactivation

A best practice for a deactivated edge CNAME configuration is to perform one of the following actions via your DNS service provider:

- Update the corresponding CNAME record to point away from our CDN service.
- Remove the corresponding CNAME record from the DNS zone.

This best practice is designed to reduce your risk exposure.

Edge CNAME Deactivation

An edge CNAME configuration may be deactivated due to any of the following actions:

- Edge CNAME configuration deletion
- Customer account deactivation

A customer account will be deactivated under any of the following conditions:

- A customer account was deleted.
- A suspended customer account experiences an extended period of inactivity.
- A trial account has expired.

Securing Your Content

Overview

There are three security tools provided to protect your content when handling HTTP and HTTPS requests. These tools are:

- **HTTP Secure (HTTPS):** This security mechanism uses Secure Socket Layer (SSL) to transmit encrypted data to ensure secure end-to-end data transmission.
- **Token-Based Authentication:** This security mechanism uses encrypted token values to ensure that the client requesting an asset is properly authorized.
- **Country Filtering:** This security mechanism prevents users from certain countries from accessing your data.
- **Web Application Firewall:** This security mechanism secures a customer origin server against malicious attacks. For more information, please refer to the Web Application Firewall administration guide.

This section will explain how the first three of these security mechanisms can be implemented.

Secure HTTP Requests (HTTPS)

If your account has been upgraded to use our Certificate Provisioning System, then you may provision a new DV or OV certificate through it. Otherwise, you need to request it via your CDN account manager.

Certificate Provisioning System

A default CDN configuration allows traffic to flow over HTTP. Serving traffic over HTTPS requires a TLS certificate on our network, updating your CDN configuration, and defining a CNAME record via your DNS service provider.

Important: Delivery over HTTPS requires TLS activation and TLS certificates. Contact your CDN account manager for more information.

TLS Certificate Support

Our CDN supports the use of any certificate derived from a Certificate Authority (CA) by allowing you to bring your own certificate to our network. Additionally, you may purchase a TLS certificate that supports multiple Subject Alternative Names (SAN). We support the following levels of validation for our certificates:

- **Domain Validation (DV):** This validation level only requires proof that you control the domain(s) associated with the certificate. Use a Domain Control Validation (DCV) method (i.e., email or DNS record) to prove your control over those domains.
- **Organization Validation (OV):** In addition to domain validation, OV requires a CA representative to contact/vet your organization and check whether the applicant may use the requested domain(s).
- **Extended Validation (EV):** EV provides the highest level of validation by requiring the CA to perform a rigorous audit to establish identity assurance. Browsers use visual cues to identify sites that leverage an EV certificate. These visual cues provide additional reassurance to your clients that they are accessing a secured site.

Quick Start

Setting up HTTPS delivery involves the following steps:

1. Activate TLS and purchase the desired TLS certificate(s). Contact your CDN account manager for more information.
2. Set up the desired TLS certificate(s).
 - **New DV or OV Certificate:** Set up a new DV or OV certificate via the Certificate Provisioning System.
 - **New EV Certificate:** Contact your account manager to set up a new EV certificate.
 - **Existing TLS Certificate:** Submit a previously purchased TLS certificate by providing the intermediate certificate, public key, and the private key via the **SSL Certificate Submission** page.
3. Enable TLS 1.3 (recommended) and/or 1.2 encryption on your web server(s).

Important: A recommended best practice is to disable support for SSL/TLS versions 1.1 or older.

4. **Customer Origin Only:** Configure your customer origin configuration for either end-to-end or client to edge encryption.
5. Create an edge CNAME configuration that points to a Subject Alternative Name (SAN) defined in the certificate (e.g., common name). Repeat this step as needed.

6. Once the TLS certificate has been deployed to our network, wait 6 hours and then verify that the certificate is live.
7. Create a CNAME record via your DNS service provider that points a SAN to the certificate's target CNAME. Repeat this step as needed.

DV and OV Certificate Setup

Setting up a DV or OV certificate requires performing the following steps:

1. Submit a request for a TLS certificate.
2. Prove your control over each Subject Alternative Name (SAN) defined in the certificate via a Domain Control Validation (DCV) method (i.e., email, DNS TXT, or DNS CNAME).
3. **OV Certificate Only:** Validate your organization by following the directions provided by the Certificate Authority (CA) via phone or email.
4. Wait until the CA has validated your request. After which, your certificate will be issued and installed on our network. Additionally, a target CNAME will be generated for each delivery platform and associated with your account.

Reminder: Once a TLS certificate has been installed on our network, you will need to update your edge CNAME configuration and DNS zone before HTTPS traffic may flow through our network. If you plan on serving HTTPS traffic via your web server(s), then you will also need to update your customer origin configuration.

To set up a DV or OV certificate

1. Navigate to the Certificate Provisioning System page. Load this page by finding the **More** menu and then selecting **Certificate Provisioning** from the **Tools** menu.
2. Click **Add Certificate**. The **New TLS Certificate** wizard will appear.
3. From the **Certificate Label** option, type a unique name that will be assigned to this request for a TLS certificate.
4. Skip the **Common Name** field. This field is auto-populated by the **Subject Alternative Name (SAN)** option.
5. From the **Domain Control Validation (DCV) Method** option, choose how you will prove your control over each domain associated with this request.
6. From the **Subject Alternative Name (SAN)** option, click **Add Domain** and then type the desired SAN. Repeat this step as needed.
7. If you have specified multiple SANs, then verify the common name defined in the **Common Name** option. Assign a different common name by clicking the **Is Common Name?** option next to the desired domain.
8. From the **Validation Level** option, select the level of validation that will be performed by the CA.

- **Domain Validation (DV):** Proceed to the next step.
- **Organization Validation (OV):** Provide your contact and organization information in the form that appears in this section. The CA will validate your organization using the provided information.

Tip: Fill out this form with care. Misspellings or incorrect information (e.g., unregistered abbreviations) will delay your request until those issues are corrected.

9. Click **Submit Request** to submit your certificate request to the CA.
10. Wait approximately 15 minutes for the CA to process your request.
11. Open your certificate request by clicking on it from the **Certificate Provisioning System** page. Verify that you are now on step **3. Domain Validation (DCV)**.
12. Validate your domain according to the DCV selected in the **Domain Control Validation (DCV) Method** option.
 - **Email:** Follow the directions for both sets of emails sent by the CA to:
 - i. The registered owner(s) of the public domain as determined by the domain's WHOIS record.
 - ii. admin@Domain, administrator@Domain, webmaster@Domain, hostmaster@Domain, and postmaster@Domain.
 - **DNS Text:** Perform the following steps:
 - i. Click **Generate** for each SAN defined in your certificate request.
 - ii. Create a TXT record for each SAN defined in your certificate request and set it to that SAN's token.
 - **DNS CNAME:** Perform the following steps:
 - i. Click **Generate** for each SAN defined in your certificate request.
 - ii. Create a CNAME record for each SAN defined in your certificate request with the following properties:
 - **Name:** Set it to a SAN-specific token.
 - **Value:** Set it to dcv.digicert.com.
13. Once the CA is able to validate your control over the domains defined in the certificate request, your request will proceed to step **4. Other Validations**. In this step, the CA performs additional validations including Organization Validation (OV).
14. **Organization Validation (OV) Only:** In addition to other validations performed by the CA, you will be required to follow the directions provided by DigiCert (CA) via an email to the contact email.

15. Once the CA has approved your request, they will issue your certificate. Our CDN service will then install it on our network and generate a target CNAME for each delivery platform.
16. Wait 6 hours and then verify that the certificate is live.
 - i. Dig the certificate's target CNAME. Note the IP address returned by the dig tool.
 - ii. Update your hosts file to point the above IP address to a SAN (e.g., common name).
 - iii. Point your browser to:
`https://SAN/`
 - iv. View the certificate by clicking on the certificate icon that appears in the browser's address bar. Verify that the common name matches the one that you requested.

Note: If the browser returns an error, then the certificate has not been fully deployed. Please wait a few more minutes and then retry.

Customer Origin Configuration

Note: This section only applies if you would like to serve HTTPS traffic from your own web servers. This section is inapplicable when setting up HTTPS delivery for CDN storage or Azure block blob storage.


A customer origin configuration determines the scheme for communication between:

- The client and the edge of our network.
- The edge of our network and your web server(s).

A prerequisite for delivery over HTTPS requires your customer origin to use the HTTPS scheme when serving content to your clients. This is known as client to edge encryption.



You may configure your customer origin to also use the HTTPS scheme when communicating with your web server(s). This is known as end-to-end encryption.




Important: Setting up HTTPS support on your customer origin configuration varies according to whether you are using customer origin groups. Customer origin groups is a new capability. Setup instructions for both customer origin groups and legacy customer origin configurations are provided below.


To update a customer origin group to use the HTTPS scheme

1. Navigate to the **Origins** page. Load this page by finding the **ADN** menu and then selecting **Customer Origin**.
2. Click on the desired customer origin group to expand it.

Note: Content delivery over HTTPS requires an edge CNAME configuration that points to this customer origin group.

3. Update origin entries to use the HTTPS scheme by performing the following steps:
 - i. From the **HTTP - Edge Protocol** section, click  next to the desired origin entry.
 - ii. From the **Protocol Type** option, set it to **HTTPS Only**.
 - iii. Click **Save**.
 - iv. Repeat the above steps as needed.

Note: The above configuration configures your origin entries to use end-to-end encryption. Alternatively, you may configure an origin entry to only encrypt traffic between the client and our network or to only encrypt traffic when the client submits a HTTPS request.

4. Review your customer origin group's TLS configuration by performing the following steps:
 - i. Click  next to the desired origin entry.
 - ii. From the **Group Settings** section, click **Origin TLS**.
 - iii. Review and adjust how our edge servers perform TLS verification with your origin servers.
 - iv. Click **Save**.

To update your legacy customer origin configuration to use the HTTPS scheme

1. Open the desired customer origin configuration.

Note: Content delivery over HTTPS requires an edge CNAME configuration that points to this customer origin.

2. Optional. Disable the use of the HTTP scheme for client to edge communication by clearing the **HTTP Edge Protocol** option.
3. Enable the **HTTPS Edge Protocol** option.
4. Configure this option using one of the following methods:
 - **End-to-End Encryption:** Specify one or more HTTPS hostnames/IP addresses under the **HTTPS Edge Protocol** option.
Sample Configuration:
https://video.mydomain.com
https://101.10.20.30
 - **Client to Edge Encryption:** Specify HTTP hostnames/IP addresses under the **HTTPS Edge Protocol** option.
Sample Configuration:
http://video.mydomain.com
http://101.10.20.30
5. Click **Update** to save your changes.

Edge CNAME Configuration

Create or update an edge CNAME configuration for each Subject Alternative Name (SAN) defined in the certificate (e.g., common name).

To set up an edge CNAME for use with a TLS certificate

1. Navigate to the **Edge CNAMEs** page corresponding to the desired platform.
2. In the **New Edge Cname** option, type the desired SAN.

Reminder: This hostname should be specified in lower-case letters and should not include a protocol (i.e., http://).

3. In the **Points To** option, select whether the edge CNAME will point to a customer origin or CDN origin server.
4. In the **Origin Directory** option, select one of the following:
 - **CDN Origin:** The selected origin directory (i.e., /00) points to CDN storage.
 - **Customer Origin:** Select the desired customer origin configuration (/80/customerorigin).

Reminder: If you point it to a customer origin, then please verify that it has been updated to leverage HTTPS.

5. Click **Add**.

Reminder: Traffic may only be delivered via this edge CNAME configuration once your DNS has been updated. Delivery over HTTPS requires a CNAME record that points a SAN to the certificate's target CNAME.

Legacy HTTPS

By default, CDN traffic may flow over HTTP. Adding support for HTTP Secure (HTTPS) requires that your CDN account manager perform the following steps on your behalf:

- The SSL traffic feature must be activated on the desired platform.
- An SSL certificate that identifies the hostname over which secure CDN traffic will flow must be installed on our network.

Request support for either of the following types of certificates:

- **Hosted SAN SSL Certificate:** This is a shared SAN certificate that has an Edgecast-owned common name. The Subject Alternative Names associated with this certificate defines the domains for several customers.
- **Custom SSL Certificate:** This is an umbrella term for any type of certificate that is dedicated to your organization. This type of certificate can define a single domain,

multiple domains, or a wildcard domain (validates unlimited subdomains). Finally, Extended Validation (EV) status can be granted to this type of certificate. An EV SSL certificate provides additional visual reassurance to your clients that they are accessing a secured site.

Note: A previously purchased SSL certificate can be leveraged for CDN usage. This process will require that we install your SSL certificate on our servers. For additional details and pricing information, please contact your CDN account manager.

Setting up HTTPS Support

Support for HTTPS traffic requires that your CDN account meet all of the following requirements:

- The SSL traffic feature for this platform must be enabled on your account.
- An SSL certificate that references the desired domain must be installed on our network.
- An edge CNAME should point the hostname referenced in the SSL certificate to a CDN or a customer origin server.

Upon requesting HTTPS support, your CDN account manager will enable the SSL Traffic feature for the platform in question and request its deployment from CDN personnel. Perform the following steps while awaiting SSL certificate activation:

1. Create an edge CNAME named after the hostname defined for the requested SSL certificate.
2. If the above edge CNAME points to a customer origin configuration, make sure that the **HTTPS Edge Protocol** option has been enabled on that customer origin configuration.
 - **End-to-End Encryption:** Set up secure communication from a client to our network to your customer origin server by specifying one or more HTTPS hostnames/IP addresses (e.g., `https://video.mydomain.com`) under the **HTTPS Edge Protocol** option.
 - **Client to Edge Encryption:** Our network can be configured to communicate over HTTPS with your clients, while transmitting data over HTTP between our network and your customer origin server. This can be accomplished by specifying HTTP hostnames/IP addresses (e.g., `http://video.mydomain.com`) under the **HTTPS Edge Protocol** option.
3. Create a CNAME on the DNS server that points the edge CNAME's hostname to the hostname provided by your CDN account manager.

Reminder: Edge CNAME configuration changes typically take effect in 1 hour.

Token-Based Authentication

Token-Based Authentication provides security for assets accessed through our content delivery network. This feature allows you to protect your content by the country, URL, IP address, protocol, or the referrer that linked to your asset. Additionally, you can protect your content by only allowing it to be available for a certain amount of time. Regardless of how you decide to protect your content, only authorized users that provide the appropriate token for the requested asset will be able to access your content. For detailed information on how you can configure Token-Based Authentication, please refer to the **Token-Based Authentication Administration Guide**, which can be downloaded either from the MCC home page or the **Token Auth** page.

Country Filtering

The default configuration for the ADN platform allows your content to be accessed from all countries. However, you can prevent users from certain countries from accessing your content. One way of accomplishing this is by configuring Token-Based Authentication to secure content by country. This type of configuration is discussed in the **Token-Based Authentication Administration Guide**. An easier implementation of authenticating by country is to use the Country Filtering feature. This section will explain how to configure this platform to only allow or block users by country.

Note: It may take up to an hour for changes to your country filtering configuration to take effect.

Administering Country Filtering Configurations

When configuring a country filter, you must specify the relative path to the location to which users will be allowed or denied access. The starting point for this relative path is defined below:

URL Type	Relative Path (Starting Point)
CDN URL	<p>Specify a relative path that starts directly after the content access point (e.g., /000001, /200001, or /800001).</p> <p>Sample URL:</p> <p><code>http://adn.0001.edgecastcdn.net/800001/customerorigin/videos/fly.flv</code></p> <p>In the above sample URL, the gray text indicates what should be excluded when securing a location. This sample request can be secured by any of the following configurations:</p> <ul style="list-style-type: none">• /• /customerorigin• /customerorigin/videos
Edge CNAME URL (CDN Origin)	<p>Specify a relative path that starts directly after the domain.</p> <p>Sample URL:</p> <p><code>http://www.domain.com/presentations/sales/businessplan.ppt</code></p> <p>In the above sample URL, the gray text indicates what should be excluded when securing a location. This sample request can be secured by any of the following configurations:</p> <ul style="list-style-type: none">• /• /presentations• /presentations/sales
Edge CNAME URL (Customer Origin)	<p>Specify a relative path that starts with the name of the customer origin configuration referenced by the edge CNAME URL.</p> <p>Sample edge CNAME URL:</p> <p><code>http://www.domain.com/Photos/Store.jpg</code></p> <p>Our edge servers will re-write the edge CNAME URL requested by the client (above) with the following CDN URL:</p> <p><code>http://adn.0001.edgecastcdn.net/800001/customerorigin/Photos/Store.jpg</code></p> <p>The above CDN URL is used to determine the starting point for the relative path that should be secured. The gray text in the above CDN URL indicates what should be excluded when securing a location. This sample request can be secured by any of the following configurations:</p> <ul style="list-style-type: none">• /• /customerorigin• /customerorigin/Photos

Key information:

- The specified relative path is compared against all CDN and edge CNAME URLs that request content from this platform.
- The path to a protected folder always starts with a forward slash (/).
- The path to a protected folder may include a trailing forward slash (/).
- It may take up to an hour before a new location is fully protected.
- This feature does not support wildcard characters (e.g., *).
- The country filtering configuration associated with the relative path will be applied recursively to that path.
- Only a single country filter can be applied to the same relative path. In other words, you cannot create multiple country filters that point to the same relative path. However, a folder may have multiple country filters. This is due to the recursive nature of country filters. In other words, a subfolder of a previously configured folder can be assigned a different country filter.

Once a relative path has been specified, choose whether to allow or block users from the specified countries.


- **Allow:** Only users from the specified countries will be allowed access to assets requested from that recursive path.
- **Deny:** Users from the specified countries will be denied access to assets requested from that recursive path. If no other country filtering options have been configured for that location, then all other users will be allowed access.

Important: The "EU" and "AP" country codes do not allow or block all requests from those regions. Rather, those country codes are applied to requests that originate from IP addresses that are spread out over a region instead of a country. If you would like to allow or block all requests from a particular region, then you would need to select all countries in that region and the region-specific code (i.e., EU or AP).


To configure country filtering options

1. Navigate to the **Country Filtering** page by finding the **ADN** menu and then selecting **Country Filtering**.
2. Click **Add Filter** to launch the **Country Filtering** wizard.
3. Specify the path to the folder (e.g., /My_Folder) to which country filtering options will be applied. Keep in mind that this path applies to both CDN and customer origin servers.
4. Select whether to **Block** or **Allow** user access to the folder specified in the previous step.
5. Type the name of each country for which access will be allowed or denied and the press ENTER.
6. Repeat the previous step as needed.
7. Click **Save**.

To modify country filtering options

1. Navigate to the **Country Filtering** page by finding the **ADN** menu and then selecting **Country Filtering**.
2. Click the  next to the desired country filtering configuration. The **Country Filtering** wizard will display the settings for the selected country filter configuration.
3. Modify the desired settings.
4. Click **Save**.

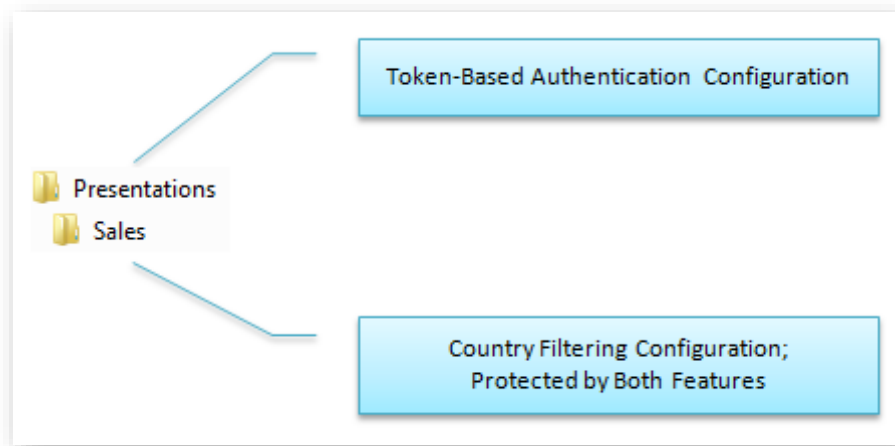
To delete country filtering options

1. Navigate to the **Country Filtering** page by finding the **ADN** menu and then selecting **Country Filtering**.
2. Click the  next to the desired country filtering configuration.
3. Click **Delete**.
4. Click **I Understand, please delete this filter**.

Protecting Assets Using Both Country Filtering and Token-Based Authentication

A Country Filtering and a Token-Based Authentication configuration cannot be applied to the same folder. However, since both of these features are applied recursively, you may apply protection on a hierarchical basis. This can be seen in the following illustration.

Note: This same principle can be used to apply multiple Country Filtering configurations to a particular folder.



In the above example, a Token-Based Authentication configuration has been applied to a folder called "Presentations." The "Presentations" folder contains a folder called "Sales." A Country Filtering configuration has been applied to this subfolder. The "Presentations" folder will be protected by Token-Based Authentication, while the "Sales" folder will be protected by both Token-Based Authentication and Country Filtering.

An asset residing in a folder that is protected by both Token-Based Authentication and Country Filtering cannot be downloaded unless the requirements for both security mechanisms are met. In the above example, content stored in the "Sales" folder cannot be downloaded unless a valid token is provided and the Country Filtering configuration does not prevent users in your country from accessing it.

Cache Management

Default Cache Management

There are three cache layers for assets routed using our ADN service, which are ADN Gateway, edge server, and user agent caching. A brief explanation for each one is provided below.

- ADN Gateway caching determines whether the requested asset will be cached on an ADN Gateway server and for how long. This type of caching determines when an ADN Gateway server will check for a new version from the origin server. By default, this caching layer is turned off.
- Edge server caching determines whether the requested asset will be cached for a particular POP and for how long. This type of caching determines when an edge server on that POP will check for a new version on the ADN Gateway server. By default, this caching layer is turned off.
- User agent caching determines whether the requested asset will be cached on a web browser and for how long. By default, this type of caching determines when the user agent will revalidate the cached asset with the edge server.

Note: Keep in mind that clearing a web browser's cache will override user agent caching.

As mentioned above, the caching mechanism for the ADN Gateway and edge servers is disabled by default. However, if your customer origin server assigns a TTL to a request, then our servers will forward the corresponding header information (i.e., Cache-Control and Expires) to the user agent that requested it.

Note: HTTP Rules Engine, which must be purchased separately, can override the above behavior. For example, it can be used to allow caching for certain types of requests on the ADN platform. For more information on the Bypass Cache feature, please refer to the **HTTP Rules Engine Administration Guide**.

Note: For detailed information on how caching works, please refer to the **How Requests are Handled** section in the **Application Delivery Network** chapter.

If you enable CDN caching for certain ADN requests, then keep in mind the following:

- By default, our CDN honors the s-maxage, max-age, and Expires header assigned to an asset. An asset's TTL is determined by those directives and the Expires header in that order. Keep in mind that s-maxage only affects the TTL for edge server caching.
- If an asset contains an Expires header that has been assigned an invalid value, including zero, then the Expires header will be ignored for that asset.
- If an asset has been assigned both a no-cache and either an s-maxage or max-age directive, then the "Cache-Control: no-cache" directive will take precedence. On a related note, if an asset has been assigned "Pragma: no-cache", then it is treated as a "Cache-Control: no-cache" directive.
- If an asset does not have header information, then it will be assigned a TTL of 7 days (i.e., Cache-Control: max-age=604800) when it is cached on our edge servers and by the user agents that requested it. You can override this default behavior by manually caching and purging assets from our POPs or through HTTP Rules Engine.
- If an edge server or an ADN Gateway server successfully revalidates a cached version of an asset, then the expiration time will automatically get updated. This means that the cached version of the asset will be assigned new Cache-Control and Expires headers based on the response returned from either the ADN Gateway server (for edge server to ADN Gateway revalidations) or the origin server (for ADN Gateway to origin server revalidations). If these headers have not been defined, then the asset will be assigned a default max-age of 7 days from the time it was successfully revalidated. Please refer to your web server's documentation for more information on how to configure these settings.

Loading Assets

By default, an asset is not cached on either the ADN Gateway or our edge servers. As a result, requests are always forwarded from our network to an origin server and then returned back to a client via our network.

As previously mentioned, the default caching behavior can be overridden through HTTP Rules Engine. If you have created a rule configuration that does so, then qualifying requests will be cached on our ADN Gateway and edge servers when it is requested by a client. This means that the first customer in a particular region to request an asset will need to retrieve the asset from either an ADN Gateway server or an origin server. If the asset's cache policy allows, then it will be cached on both servers. After which, future requests will be able to retrieve the asset directly from either the edge server or ADN Gateway server. This translates to a slightly longer load time for the first client to request a cacheable asset. If you have a large event or content that becomes simultaneously available to a large volume of end-users (e.g., a new movie release),

then it might be helpful to load your media content on an ADN Gateway server and our edge servers.

Note: There is a default limit of 50 concurrent load requests at any given time.

Note: An asset can be automatically or manually purged. For more information on how purging works, please refer to the **Purging Files** section below.

Note: Loading an asset will cause it to be cached across the CDN regardless of the CNAME used to reach it; provided that those CNAME records point to the same location.

To load content

1. Navigate to the **Purge/Load** page by finding the **ADN** menu and then selecting **Purge/Load**.
2. From the **URL** option, which can be found under the **Load To Edge** section, select the base URL where the desired content can be found.
3. Directly below the **URL** option, type the relative path to the asset that will be loaded. Since a forward slash (/) has been added for your convenience, make sure that this relative path starts after the forward slash.
 - The base URL for this relative path is the URL selected in the previous step.
 - Make sure to include the file name extension when specifying the file name.
4. Click **Load Content**.

Note: Each load request needs to be processed. Therefore, your request will be placed in a queue. The **Purge/Load** page provides a means to keep track of when a request was made and when it was completed.

To perform a bulk load

1. Navigate to the **Purge/Load** page.
2. Click **Bulk Load**, which can be found under the **Load To Edge** section, to display the **URLs** option.
3. In the **URLs** option, type the CDN or edge CNAME URL that points to the asset(s) that will be loaded.
4. Add another load request by pressing ENTER and then specifying the desired CDN or edge CNAME URL. Repeat this step as needed.
5. From the **Regions** option, perform one of the following steps:
 - **Load to All Regions (Recommended):** Select **All**.
 - **Load to Specific Regions:**
 - i. Select **Custom**.
 - ii. Mark the regions where content should be loaded.
 - iii. Clear all other regions.
6. Click **Load Content**.

Note: Each CDN or edge CNAME URL specified in the **URLs** option must be placed on a separate line. This can be accomplished by using a carriage return to delimit each CDN or edge CNAME URL. Using any other type of character (e.g., a comma) as a delimiter may prevent the asset from being loaded.

Automatically Loading Assets through the Web Services REST API

Content may be automatically loaded via our Web Services REST API. For more information, please refer to our [REST API Help Center](#).

Reminder: Although the Web Services REST API allows you to quickly submit multiple load requests, the default limit of 50 concurrent load requests is still applicable.

Purging Assets

By default, content requested through the ADN platform is not cached. As a result, it is not necessary to purge content for this platform. However, the default caching behavior can be overridden through HTTP Rules Engine. If you have created a rule configuration that does so, then qualifying requests will be cached on our ADN Gateway and edge servers. You may wish to purge this content when you would like to make a newer version of that asset available to your clients. Keep in mind that an asset will continue to be served out until the content of that asset expires or it is purged.

Important: Purging is not the same thing as deleting the asset from the origin server. Purging an asset will not delete the source asset. It will simply remove the cached instance of that asset from the appropriate ADN Gateway server and all of our edge servers. If you would like to remove the original asset, then you will need to delete it from the origin server through your preferred FTP client.

Note: There is a default limit of 50 concurrent purge requests at any given time.

Automated Purging Through the Web Services REST API

Reminder: By default, content requested through the ADN platform is not cached. This means that a cached version of an asset will typically not reside on an ADN Gateway or edge server.

Content may be automatically purged via our Web Services REST API. For more information, please refer to our [REST API Help Center](#).

Reminder: Although the Web Services REST API allows you to quickly submit multiple purge requests, the default limit of 50 concurrent purge requests is still applicable.

Manually Purging Assets

Reminder: By default, content requested through the ADN platform is not cached. This means that a cached version of an asset will typically not reside on an ADN Gateway or edge server. As a result, it is not necessary to purge content from the ADN platform.

Cached assets can be purged from our ADN Gateway and edge servers through the **Purge/Load** page.

Keep the following information in mind when purging content:

- A request to purge content is specific to the following criteria:
 - **Directory:** A URL must be specified when purging content. This URL identifies the directory where the content to be purged is stored.
 - Purging is not URL-specific. In other words, it affects all CDN and edge CNAME URLs that point to the same asset.
 - **Region:** Purging by region is an advanced option that should only be used to address specialized cache management needs. It allows the selection of one or more regions when purging content. The asset in question will only be purged from the selected region(s).
 - By default, purging is applied to the entire network (i.e., all regions).
- Purges are protocol-independent (i.e., HTTP or HTTPS).
- An asset only needs to be purged a single time per directory and region.
- Each purge request needs to be applied to edge servers throughout our network. As a result, each request will be placed in a queue. The **Purge/Load** page provides the means to keep track of when a request was made and when it was completed.
- A URL containing a dollar sign symbol (i.e., \$) cannot be added to the purge queue. The dollar sign symbol must be encoded as "%24" before such a URL can be added to the queue.
- Use purge syntax to determine the scope of the purge request. You can choose to purge assets individually, by folder, or recursively. Detailed information on how to perform each type of purge is described in the corresponding sections below.
- If you have enabled caching on the ADN platform and you have assets that are updated frequently, then you should consider using a naming convention that assigns unique names to your new assets. This will ensure that future requests will go to the new assets without the need of manually purging the old assets.
- An alternative to changing your file naming convention for assets that are updated frequently is to use the **URLs** option. This option allows you to specify a list of assets

that will be placed into the purge queue. If you keep a record of frequently purged assets, then you can copy and paste them into this option.

- There is a default limit of 50 concurrent purge requests at any given time.
- Each purge request needs to be applied to edge servers throughout our network. As a result, each request will be placed in a queue. The **Purge/Load** page provides a means to keep track of when a request was made and when it was completed.

To manually purge assets

1. Navigate to the **Purge/Load** page by finding the **ADN** menu and then selecting **Purge/Load**.
2. From the **URL** option, which can be found under the **Purge From Edge** section, select the base URL where the desired content can be found.
3. Directly below the **URL** option, type the path to the folder containing the asset(s) to be purged. Since a forward slash (/) has been added for your convenience, make sure that this relative path starts after the forward slash.
4. Perform one of the following:

Purge	Procedure	Sample Syntax
Asset	Append the name of the asset that you would like to purge to the relative path specified in step 3.	Sales/Gala_01.flv
Set of Assets	Determine the naming convention and/or file type that will be used to identify the assets that will be purged. Append that pattern to the relative path specified in step 3.	Sales/Gala*.fl*
Folder	Append an asterisk, a period, and another asterisk (i.e., /*.*) to the relative path specified in step 3.	Sales/*.*
Folder (Recursively)	Append an asterisk (i.e., /*) to the relative path specified in step 3.	Sales/*

5. From the **Regions** option, perform one of the following steps:
 - **Purge All Regions (Recommended):** Select **All**.
 - **Purge Specific Regions:**
 - i. Select **Custom**.
 - ii. Mark the regions that should be purged.
 - iii. Clear all other regions.
6. Click **Purge Content**.

To perform a bulk purge

1. Navigate to the **Purge/Load** page by finding the **ADN** menu and then selecting **Purge/Load**.
2. Click **Bulk Purge**, which can be found under the **Purge From Edge** section, to display the **URLs** option.
3. In the **URLs** option, type the CDN or edge CNAME URL that points to the asset(s) that will be purged.
4. Add another purge request by pressing ENTER and then specifying the desired CDN or edge CNAME URL. Repeat this step as needed.
5. From the **Regions** option, perform one of the following steps:
 - **Purge All Regions (Recommended):** Select **All**.
 - **Purge Specific Regions:**
 1. Select **Custom**.
 2. Mark the regions that should be purged.
 3. Clear all other regions.
6. Click **Purge Content**.

Note: Each CDN or edge CNAME URL specified in the **URLs** option must be placed on a separate line. This can be accomplished by using a carriage return to delimit each CDN or edge CNAME URL. Using any other type of character (e.g., a comma) as a delimiter may prevent the asset from being purged.

Purging an Individual Asset

An individual asset can be purged by specifying a CDN or an edge CNAME URL that points to it. Keep in mind that this type of purge request will also purge all query string variations of the specified asset at that location. Sample URLs are provided below.

- **CDN URL:** `http://adn.0001.edgecastcdn.net/000001/resources/homepage.html`
- **Edge CNAME URL:** `http://cdn.sampleurl.net/resources/homepage.html`

Purging a Set of Assets

A pattern can be defined through the use of asterisks to determine the assets that will be purged from the specified directory. Each asterisk represents one or more characters. If you would like to purge a particular file type, then you can simply append an asterisk, a period, and the desired filename extension to the purge URL (e.g., `http://cdn.sampleurl.net/*.html`).

A list of sample purges is provided below.

Example	Description
<code>http://cdn.sampleurl.net/folder01/*.*</code>	This sample purge request will purge the entire contents of the directory called "folder01." This purge request will not be performed recursively.
<code>http://cdn.sampleurl.net/folder01/*.css</code>	This sample purge request will purge all CSS assets from the directory called "folder01."
<code>http://cdn.sampleurl.net/folder01/a*.ht*</code>	This sample purge request will purge all assets that start with the letter "a" and whose filename extension starts with "ht" (e.g., activity.html).

An asterisk (*) can either be used as a wildcard in a filename pattern or to recursively purge a folder's content. However, it cannot be used as a wildcard when specifying a directory. Additionally, you should keep in mind a recursive purge is always performed for purge requests that end with a forward slash followed by an asterisk (/*). If you would like to only purge a folder's content, append /*.* to the desired purge directory.

Recursive Purging

We have already seen how to purge the contents of a folder through the use of wildcards (i.e., *.*). If you would prefer to purge a directory and all of its subfolders, then you should specify a purge request for that directory and append an asterisk. A sample purge request is provided below.

- `http://cdn.sampleurl.net/folder/*`

Reminder: An asterisk cannot be used to specify a folder pattern. It can only be used as specified above or when defining a pattern that will be used to define the set of assets that will be purged (e.g., /*.css).

Purging Query String Variations of a Cached Asset

If your file naming structure is based on query strings, you can customize how query string variations of that asset are handled. Several examples on how you can handle query string variations of an asset are provided below.

Sample Purge Request	Description
<code>http://cdn.sampleurl.net/folder/sampleasset.htm</code>	Simply specifying a filename will purge the original asset and all of its query string variations.
<code>http://cdn.sampleurl.net/folder/sampleasset.htm?</code>	A filename followed by a question mark (?) indicates that only the specified asset will be purged. All query string variations of that asset will not be purged.
<code>http://cdn.sampleurl.net/folder/sampleasset.htm?Query</code>	A filename followed by a question mark and a query string (?SampleQuery) indicates that only the specified asset with that particular query string will be purged. The original asset and all other query string variations of that asset will not be purged.

Purging a Directory's Default Cached Asset

The default asset cached for a particular directory can be purged through the use of the syntax described below.

Sample Purge Request	Description
<code>http://cdn.sampleurl.net/folder/</code>	The default cached asset for a directory, along with all of its query string variations, can be purged by specifying a CDN or edge CNAME URL to the desired directory.
<code>http://cdn.sampleurl.net/folder/?</code>	The default cached asset for a directory, excluding query string variations, can be purged by specifying a CDN or edge CNAME URL to the desired directory and appending a question mark.
<code>http://cdn.sampleurl.net/folder/?Query</code>	A query string variation of the default cached asset can be purged by specifying a CDN or edge CNAME URL to the desired directory and appending the desired query string.

Note: When purging the default asset that was cached for a particular folder, keep in mind that there are certain conditions that may cause an asset other than the default asset for that folder to be purged. For example, specifying "/Folder/" may not be the same as specifying "/folder/index.html" even if "index.html" is the default asset for that folder. A couple of situations under which this would occur are when the specified URL is redirected or if URL rewriting is used.

Purge and Load History

A log of the most recent purge and load requests is provided on the **Purge/Load** page. Purge and load history indicates:

- The CDN URL to the asset that was purged/loaded (**Path**).
- The set of regions to which the purge/load request was applied.
- When the request was made (**Created**).
- The date/time when the action was completed (**Completed**). If the **Completed** column is blank, then the request is still being processed.

Tip: Purge/load history is paginated. Leverage the navigation bar at the bottom left of the page to navigate between pages.

Note: Purge and load history is only displayed for the last 200 requests within the last 7 days. Requests older than a week will not be displayed. Additionally, once this limit has been reached, old purge or load requests will be cleared to make way for new requests.

Note: Purge/load history is reported using the GMT time zone.

Query String Caching

Reminder: By default, content requested through the ADN platform is not cached. The options discussed in this section only affect cacheable content.

Requests that contain a query string can be cached differently. You can choose to cache these types of assets in one of three different ways. Each cache method is explained below.

Cache Method	Description
Standard-cache	This is the default mode for handling URLs that contain query strings (e.g., Asset.txt?Data=True). This mode will ignore query strings in the URL. It will only cache a single asset per URL, regardless of the presence of query strings. However, if there is a cache miss or an expired cache asset needs to be revalidated, then the full CDN URL, including query strings, will be passed to the customer origin server.
No-cache	This mode prevents requests containing query strings from being cached. All requests that contain query strings will be forwarded to the customer origin server.
Unique-cache	This mode will cache an asset for each request made with a unique URL. A unique cache asset will be created by even the slightest variation in the query string. It is recommended that you enable the query string logging feature when taking advantage of unique-cache mode. Additionally, it is highly recommended that this mode be avoided when requests will contain query strings that will result in a very low cache hit ratio. For example, information such as session ID or user name would create such a scenario.

Token-Based Authentication and Query String Caching

By default, no-cache or unique-cache mode cannot be used with Token-Based Authentication. This means that if you have protected at least one folder with Token-Based Authentication, then the only cache method that will be available from the MCC is standard-cache. Likewise, if you have configured query string caching to use no-cache or unique-cache mode, then you will not be allowed to protect a folder with Token-Based Authentication.

If you would like to protect your assets using Token-Based Authentication and you have other assets that require query string caching, then a custom configuration will have to be implemented for your account. This custom configuration will disable query string caching on folders that have been protected by Token-Based Authentication. If you would like such a custom configuration, then you should contact your CDN account manager.

Note: A custom configuration, such as the one described above, will cause the options that appear on the **Query-String Caching** page to be disabled.

Query String Logging

Query string logging determines whether the query string portion of a CDN or edge CNAME URL will be recorded in our raw access logs. Raw access logs are the log files from which we generate our reports. Keeping track of query string information is important when you have enabled unique-caching mode, which caches assets according to query string parameters. By keeping track of query string information, you can view statistics on each cached asset.

Asset Compression

Compression allows you to save on bandwidth costs and reduces the amount of time that it takes for your users to download your content. An asset can be compressed either by an origin server or an edge server. Both compression options are described below.

Customer Origin Server Compression

An origin server may compress an asset before delivering it to an edge server.

Setup

Origin server compression requires the following setup:

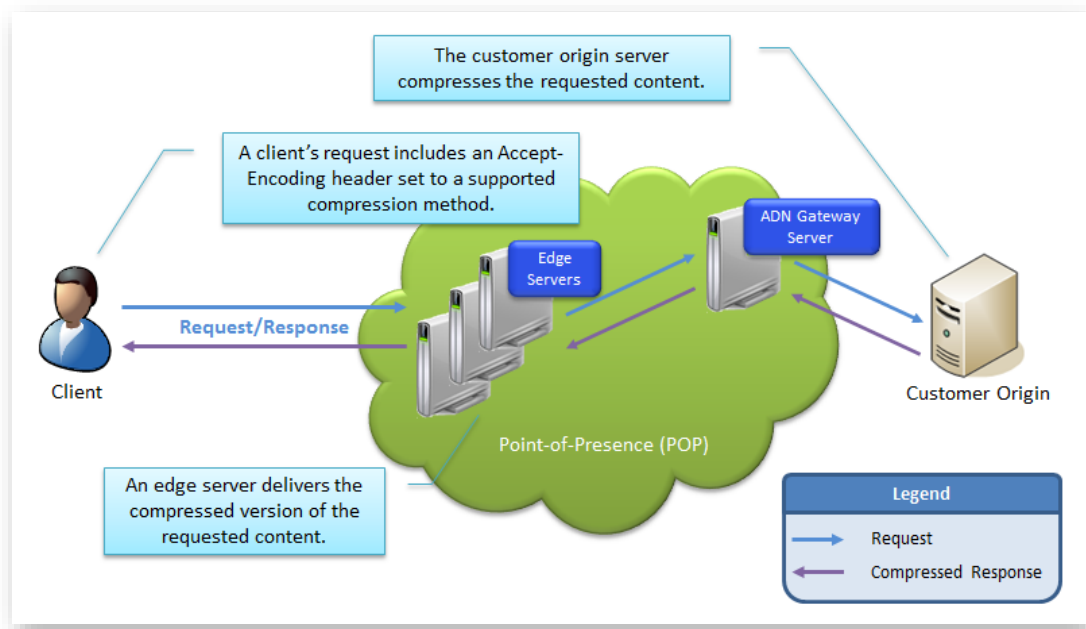
1. The request must include an Accept-Encoding header set to one or more of the following compression types:
 - gzip
 - deflate
 - bzip2
 - br

Note: This header allows the user agent to indicate which compression methods it supports to the origin server.

2. The web servers associated with your customer origin must support the compression method defined within the Accept-Encoding request header.

Basic Workflow

The process through which requested content is compressed is illustrated below.



HTTP Request/Response (Origin Server Compression)

Detailed Workflow

Detailed information on how requests are handled with regards to origin server compression is provided below.

1. **Requester:** A user submits a request for content. The manner in which this request is handled by our edge servers depends on whether the request includes the Accept-Encoding header.

Accept-Encoding Header	Description
Set to a supported compression method.	This type of request will be handled as described in step 2.
Set to an unsupported compression method.	This type of request is uncacheable. Our edge servers will always retrieve the requested content from your customer origin server.
Missing	This type of request will be served in an uncompressed format.

2. **POP:** An edge server on the POP closest to the client will check to see if the requested content has been cached and if it still has a valid TTL.
- **Cache Miss:** If a cached version of the requested content is not found, then the request will be forwarded to an origin server. Proceed to step 3.
 - **Cache Hit & Matching Compression Method:** An edge server will immediately deliver the compressed content to the client.
 - **Cache Hit & Different Compression Method:** If the client requests a supported compression method that is different from the one used by the initial request, then an edge server will transcode the asset to the requested compression method.
 - **Uncompressed Cache Hit:** If the initial request caused the asset to be cached in an uncompressed format, then a check will be performed to see whether the request is eligible for edge server compression.
 - **Eligible:** The requested asset will be compressed as indicated in the **Edge Server Compression** section and then delivered to the client. Your caching policy dictates whether the compressed asset is eligible to be cached on both the respective edge server and the ADN Gateway server.
 - **Ineligible:** An edge server will immediately deliver the uncompressed content to the client.

Note: By default, content requested through the ADN platform will not be cached and therefore should result in a cache miss.

3. **Origin Server:** The manner in which the origin server will handle the request depends on whether both the CDN and your web servers support the compression method specified in the Accept-Encoding header.

CDN	Customer Origin	Request Workflow
Supported compression method	Supported compression method	<ol style="list-style-type: none"> 1. Customer Origin: It will compress the asset and send it to an edge server via an ADN Gateway server. 2. Edge Server: It will serve the compressed asset to the client. Your caching policy dictates whether the compressed asset is eligible to be cached on both the ADN Gateway server and edge server.

CDN	Customer Origin	Request Workflow
Supported compression method	Unsupported compression method	<ol style="list-style-type: none"> 1. Customer Origin: It will serve an uncompressed asset to an edge server via an ADN Gateway server. 2. Edge Server: It will: <ul style="list-style-type: none"> • Serve the uncompressed asset to the client. • Check to see whether the request is eligible for edge server compression. <ol style="list-style-type: none"> i. Eligible: The requested asset will be compressed as indicated in the Edge Server Compression section and then delivered to the client. Your caching policy dictates whether the compressed asset is eligible to be cached on both the ADN Gateway server and edge server. ii. Ineligible: An edge server will immediately deliver the uncompressed content to the client. Your caching policy dictates whether the uncompressed asset is eligible to be cached on both the ADN Gateway server and edge server.
Unsupported compression method	Supported compression method	<ol style="list-style-type: none"> 1. Customer Origin: It will serve a compressed asset to an edge server via an ADN Gateway server. 2. Edge Server: It will serve the compressed asset to the client. However, it will not cache it.
Unsupported compression method	Unsupported compression method	<ol style="list-style-type: none"> 1. Customer Origin: It will serve an uncompressed asset to an edge server via an ADN Gateway server. 2. Edge Server: It will serve the uncompressed asset to the client. However, it will not cache it.

Edge Server Compression

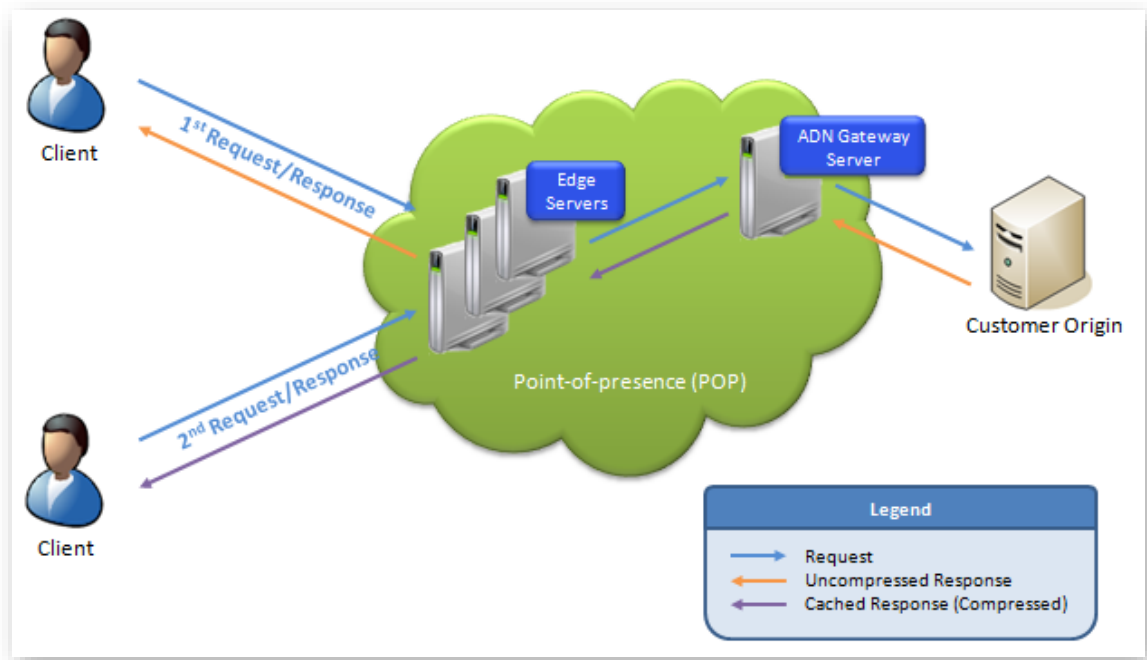
Reminder: By default, content requested through the ADN platform is not cached. Edge server compression can only be applied to cacheable content.

Our edge servers can be configured to compress certain types of assets before they are delivered to your clients. The type of assets that will be compressed is determined by a list of content types that have been specified in the **File Types** option of the **Compression** page. Content type identifies an asset by type and subtype. For example, a content type of "text/plain" identifies text assets, while a content type of "text/html" identifies HTML assets.

Key information:

- Edge server compression requires the Accept-Encoding header and it must be set to a supported compression method. Supported compression methods are:
 - gzip
 - deflate
 - bzip2
- Content compressed by our edge servers will only be served from cache. As a result, edge server compression is only suitable for cacheable content.
- The **File Types** option accepts one or more content types that will be compressed.
- Multiple content types can be specified by separating each one with a comma.
- Compression can only be applied to content smaller than 3 MB. Larger assets will not be compressed by our servers.
- Certain types of content, such as images, video, and audio media assets (e.g., JPG, MP3, MP4, etc.), are already compressed. Additional compression on these types of assets may not significantly diminish file size. Therefore, it is recommended that you do not enable compression on these types of assets.
- Wildcard characters, such as asterisks, are not supported in the **File Types** option.

The process through which our edge servers compress content is illustrated below.



HTTP Requests/Responses (Edge Server Compression)

The process through which requested content is compressed is outlined below. This process assumes that a client performed a request that is eligible for edge server compression.

1. An edge server on the POP closest to the client will check to see if the requested content has been cached and if it still has a valid TTL.
 - **Compressed Cache Hit:** An edge server will immediately deliver the compressed content to the client.
 - **Uncompressed Cache Hit:** An edge server will compress the requested content, serve it to the client, and then cache the compressed asset.
 - **Cache Miss:** If the asset has not been cached, then the request will be forwarded to an origin server. Proceed to the next step.
2. The response from the origin server will be one of the following:
 - **Compressed Content:** An edge server will deliver the compressed content to the client. For more information, please refer to the **Origin Server Compression** section above.
 - **Uncompressed Content:** An edge server will serve the uncompressed content to the client. After which, the edge server will compress it and then cache the compressed version.

To set compression options

1. Navigate to the **Compression** page by finding the **ADN** menu, pointing to **Cache Settings**, and then selecting **Compression**.
2. Make sure that the **Compression Enabled** option is selected.
3. In the **File Types** list, make sure that only the desired content types are listed. For each content type that is not currently listed, append a comma and the content type (i.e., type/subtype).
4. Click **Update** to save your changes.

Note: Keep in mind that it may take up to an hour for your changes to take effect.

Overwriting Cache Headers

Another way to override the default cache behavior is to overwrite the cache headers assigned to your assets. This can be accomplished through HTTP Rules Engine. This feature allows you to create custom rules that can be used to handle how our ADN Gateway and edge servers cache and grant access to assets requested through HTTP-based platforms. For detailed information on how to use HTTP Rules Engine, please refer to the **HTTP Rules Engine Administration Guide**, which is available from the MCC home page and the **Rules Engine** page.

Note: HTTP Rules Engine must be purchased separately. For more information, please contact your CDN account manager.

Glossary

A

Asset

This term refers to a resource that contains header information and a body that can be served to clients. Examples of assets include files and dynamic content.

C

Cache

This term refers to the storage of data to improve data delivery performance. When used in reference to our CDN, it refers to the temporary storage of an asset on an edge server. Cache increases the speed through which that particular edge server can deliver that asset for subsequent requests.

CDN

Our content delivery network (CDN) consists of points-of-presence (POPs) that are placed at critical network and geographical locations around the world. This allows us to place content at the edge of the Internet allowing for faster downloads by your end-users.

CDN Domain

This term refers to a domain name assigned to your account. In the following examples of CDN domains, *xxxx* represents your CDN account number.

- `adn.xxxx.edgecastcdn.net`
- `wac.xxxx.edgecastcdn.net`
- `wpc.xxxx.edgecastcdn.net`
- `fms.xxxx.edgecastcdn.net`

CDN Origin

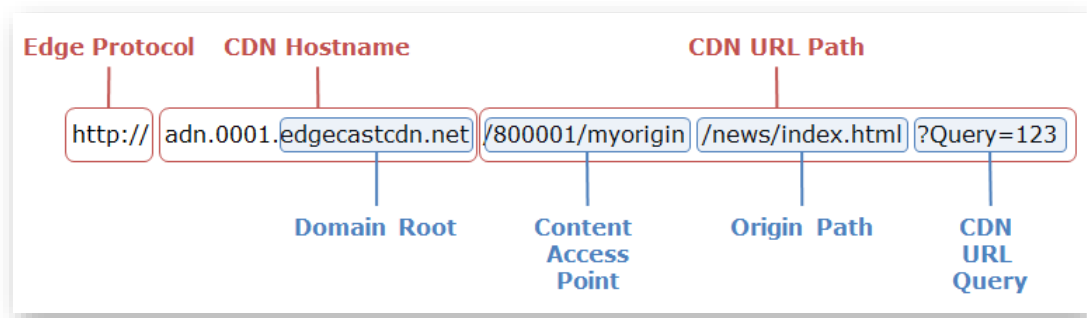
This term refers to a storage server on our CDN. Our CDN origin servers are in close proximity to our POPs, in order to provide optimal conditions for transferring data from a CDN origin server to your end-users via our POPs.

CDN Origin Identifier

This type of identifier in the CDN URL indicates that requested asset should be retrieved from the CDN origin server. A CDN origin identifier is indicated by "00" as the starting two numbers in the CDN URL path.

CDN URL

This type of URL identifies a location or an asset on our content delivery network. The following diagram indicates the different components in a CDN URL. Keep in mind that xxxx represents your CDN account number.



CDN URL Path

This term refers to the portion of the CDN URL that appears after the CDN domain. It provides the relative path to a folder or an asset on either a CDN or customer origin server. In the following examples of CDN URL paths, xxxx represents your CDN account number.

- `/00xxxx`
- `/00xxxx/Videos/Sales/11/`
- `/00xxxx/Videos/Sales/11/Presentation01.flv`

CDN/Edge CNAME URL Query

This term refers to the query string that appears after a question mark in a CDN or edge CNAME URL. If Token-Based Authentication is protecting the requested content, then a token value should appear directly after the question mark.

CNAME

A Canonical Name (CNAME) record is used to indicate that a domain name is an alias of another domain name. A CNAME record must be registered on a Domain Name System (DNS). This term should not be confused with edge CNAME.

Content Access Point

It provides a point of reference to any folder on a CDN or customer origin server. This relative path starts directly after the CDN domain. The proper syntax for a content access point is `"/yyxxx/path,"` where *yy* stands for the identifier and *xxx* stands for the CDN account number. The term *path* is optional and stands for the path to the folder specified by an edge CNAME configuration.

Customer Origin

This term refers to a storage server that is external to our CDN. Assets can be delivered from your storage server to your end-users via our POPs.

Customer Origin Identifier

This type of identifier in the CDN URL indicates that requested asset should be retrieved from the customer origin server. A customer origin identifier is indicated by "80" as the starting two numbers in the CDN URL path.

D

Domain Root

This term identifies the top and second-level domains associated with the CDN domain name. An example of a domain root is "google.com."

Dynamic Streaming

This technology, also known as Adaptive Streaming, allows a player (e.g., Silverlight) to dynamically switch between bit rate streams of varying quality levels, in order to provide an optimal viewing experience based on a client's bandwidth and CPU usage. Smooth Streaming is an example of adaptive streaming.

E

Edge CNAME

This term refers to the mapping of a CNAME record to a directory on a CDN or customer origin server. The purpose of this mapping, which is only used by our CDN, is to establish a user-friendly alias for content served through the CDN. It relies upon your CNAME record being properly mapped on a DNS server.

Edge CNAME URL

This type of URL takes advantage of an edge CNAME to mask a CDN URL. This allows it to identify a location or an asset on our content delivery network using a more user-friendly URL. An edge CNAME URL is specific to the platform from which it was configured.

In the following examples, the domain assigned to the edge CNAME is "www.mydomain.com."

In the first example, the edge CNAME references the following CDN URL:

"http://adn.xxxx.edgecastcdn.net/00xxxx." In the following two examples, the edge CNAME references the following CDN URL: "http://adn.xxxx.edgecastcdn.net/00xxxx/Videos."

Edge CNAME URL	Points To
http://www.MyDomain.com/	http://adn.xxxx.edgecastcdn.net/00xxxx/
http://www.MyDomain.com/Sales/11/	http://adn.xxxx.edgecastcdn.net/00xxxx/Videos/Sales/11/
http://www.MyDomain.com/Sales/11/Presentation01.flv	http://adn.xxxx.edgecastcdn.net/00xxxx/Videos/Sales/11/Presentation01.flv

Edge CNAME URL Path

This term refers to the portion of the edge CNAME URL that appears after the edge CNAME. It provides the relative path to a folder or an asset on a CDN or customer origin server. In the following examples of edge CNAME URL paths, the edge CNAME points to the following CDN URL: "http://adn.xxxx.edgecastcdn.net/00xxxx/Videos."

Edge CNAME URL Path	Actual Edge CNAME URL
/Sales/11/	http://adn.xxxx.edgecastcdn.net/00xxxx/Videos/Sales/11/
/Sales/11/Show01.flv	http://adn.xxxx.edgecastcdn.net/00xxxx/Videos/Sales/11/Show01.flv

Edge Protocol

This term refers to the protocol (e.g., HTTP and RTMP) used in a CDN URL or an edge CNAME URL.

Edge Server

This type of server is located near the edge of the Internet where its close proximity to your end-users allows it to deliver data more quickly than normal Internet communications. Our edge servers are integral component of our POPs.

Encryption Key

Token-Based Authentication requires the use of an encryption key to encrypt and decrypt token values. There are two types of encryption keys, which are a primary and a backup key. Both of these keys can be used to encrypt and decrypt token values.

G

Global Key

This type of Live Authentication key can be used to authenticate live streams. Only a single global key can be specified.

H

HTTP Large

This platform consists of dedicated edge servers that retrieve, cache, and serve large assets to your clients. These servers have been optimized to cache assets. A typical asset for the HTTP Large platform is larger than 300 KB.

HTTP Progressive Download

This method of streaming video content is performed through the HTTP protocol. Progressive downloads are not as secure as other streaming methods, since the entire asset will be stored on your client's computer. This allows your client to save and share your content with other users.

HTTP Small

This platform consists of dedicated edge servers that retrieve, cache, and serve smaller content to your clients. These servers have been optimized to index files. A typical asset for the HTTP Small platform is smaller than 300 KB.

I

Identifier

It identifies how a request will be routed through our CDN. Examples of identifiers are:

- **00**: CDN origin identifier
- **80**: Customer origin identifier

Load

This feature allows you to cache an asset on all of our POPs. This feature is unsupported for use with our live streaming solutions.

M

Media Control Center (MCC)

This web application is provided to help you manage all of your CDN needs. The major features that are available from the MCC are CDN configuration settings, cache management, file management, reports, and analytics. Additionally, the MCC allows you to configure your organization's settings, such as granting or denying access to the MCC. You can access the MCC through the following URL:

<https://my.edgecast.com>

O

Origin Path

It references a relative path to a folder or an asset in a CDN URL. This type of path follows the content access point.

Origin Server

This term refers to the servers that store the assets that will be distributed by our POPs. There are two types of origin servers, which are CDN origin and customer origin servers.

P

Player URL

A media player uses this type of URL to stream content. It identifies the location of the streamer on the live ingestion point.

Point-of-Presence (POP)

A point-of-presence, or data center, is an access point to the Internet. The main components of a POP are edge servers, CDN origin servers, and publishing servers.

Pre-Cached

A pre-cached asset means that it has been loaded to all of our POPs. Pre-caching your assets allows even quicker content delivery to your clients, since it ensures that the requested asset will not have to be retrieved from the origin server.

Purge

This feature allows you to remove the cached version of an asset from all of our edge servers. A purge can be performed on a folder or an individual asset.

Q

Query String

Additional data can be appended to a URL (e.g., `http://www.server.com/index.html?Data=xyz`). This information can be used in a variety of ways. Our CDN allows you to leverage this information to determine how content will be cached. Additionally, you can choose to store query string information in our log files.

R

Request

A request consists of a set of headers and a body sent from a client. This header data and the body define the requested content. Typically, a request is sent from a client to an edge server. If the requested content is not found, then our edge servers will forward this request to an origin server.

Response

A response consists of the headers and the body sent from a server responding to a request. If an origin server is returning a response, then this response will be sent to an edge server. The edge server will then forward the response to a client.

T

Time to Live (TTL)

This term refers to the amount of time that a cached asset is still considered fresh. Our edge servers will continue to serve a cached version of an asset while its TTL has not expired. An asset's TTL is calculated by the Cache-Control and Expires headers associated with the response sent by a CDN or customer origin server.

Token

A token value must be provided when Token-Based Authentication has been applied to the requested content. Each token value contains a set of encoded requirements that must be met before content delivery may take place. A token value may be specified as a query string in the request URL (e.g., `sales.pdf?1234567890AB`).

Token-Based Authentication

This capability provides the means for defining the set of content that will require authentication prior to delivery. Authentication takes place via an encoded token value that must be included in the request URL. This token value is then decrypted on an edge server. The requested content will only be delivered when the user meets the requirement(s) defined in the token.