Limelight

Performance at the Edge

Limelight Control User Guide

Limelight Control Version 22.5

2022-05

Table of Contents

Control Portal Overview

Limelight's web-based customer portal, Limelight Control, provides 24x7 access to Limelight services and support. You can use the portal to order and configure many services, manage content, analyze usage, and access online support.

Supported Browsers

Limelight Control is tested against and supports the latest version of the following browsers:

- Firefox
- Safari
- Chrome

Menus and Links

The Control Portal layout provides an optimized user experience with easy-to-access functionality. The navigation pane is displayed on the left side of the Control Portal. Icons for making other settings are located at the top right of the Control Portal.

Navigation Pane and Icons

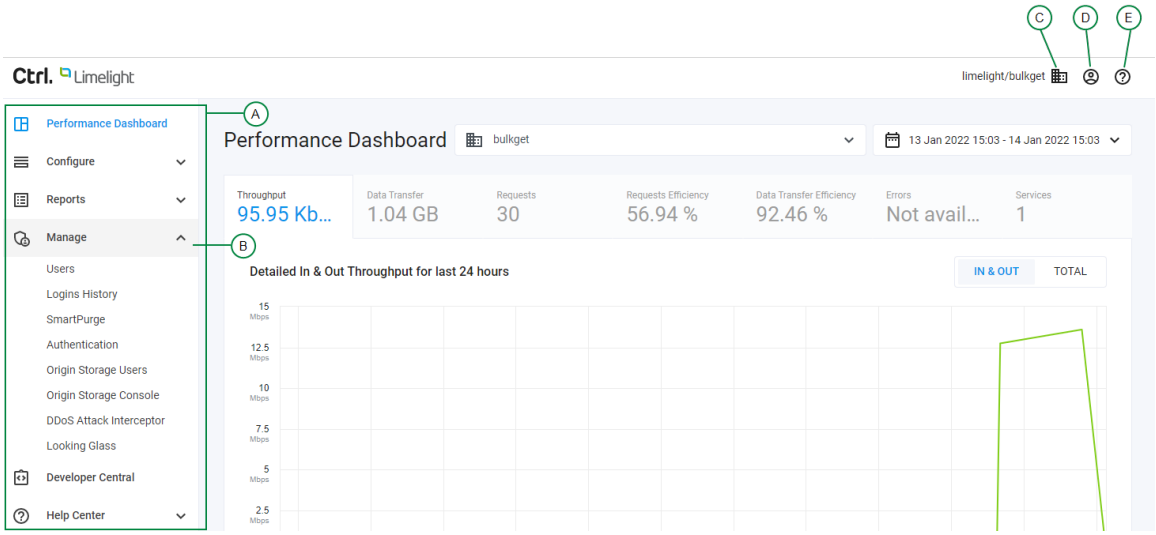

Item	Description
A	Navigation pane with expandable menus. See Menus for more information.
B	Expanded menu with menu items. Click an item to open the corresponding page.
C	Company/Account icon. Click and select the company and account you want to work with in Control. **Note:** Accounts are also known as "shortnames."
D	Profile icon. Click to search for users, edit account settings, and log out of Control.
E	Context help icon. Visible on each screen in Control. Click to view documentation specific to the screen.

Menus

Item	Description	More Information
Dashboard	View summary about your account.	Dashboard

Item	Description	More Information
Configure	Configure your products and services.	Configure
Reports	View reports on your traffic, network efficiency, and more.	Reports
Manage	Manage users, view login history, and other functionality.	Manage
Developer Center	Access a rich library of tools and resources to help developers get the most from Limelight APIs.	Developer Central
Help Center	View Limelight's network status, customer documentation, open a support ticket, or contact Limelight Support.	Click any of the menu items.

Dashboard

The *Limelight Control* Dashboard provides information about the following items:

- Information such as segments and number of configuration for each of your accounts
- Recent activities and events regarding your accounts
- Popular data segments related to your accounts

You can access the Dashboard by clicking the Dashboard (see Navigation Pane and Icons) menu in the navigation pane.

> **Note:** You can view traffic performance metrics on the Performance Dashboard.

Accounts

The Accounts section displays summary information for each Account within the currently-selected Company. The default view shows the first three Accounts. You can view other accounts by selecting one of the position markers immediately below. The marker for the currently displayed Account is shown in gray.

About Data Segments

Data Segments are filters that gather traffic data. They are powered by EdgeQuery, and are used by the Traffic Report and these endpoints:

- /realtime-reporting-api/traffic
- /realtime-reporting-api/traffic/geo

Each Data Segment gathers the following data for a source or published host, a specific account, and a specific protocol (HTTP, HTTPS):

- Total traffic throughput.
- Percentage of total traffic relative to all Data Segments.
- Number of requests.
- CDN efficiency.

Data Segments gather daily data, providing 24-hour latency, and granularity. A Data Segment can optionally gather realtime data, providing 5-minute latency and granularity.

Data segments require significant system resources to process. A Data Segment is considered dormant if the Data Segment has not collected data within the last 100 days. Segments without data are removed weekly to improve EdgeQuery backend performance.

Account Summary Information

The summary information provided for each Account includes:

- **segments** - If Segments are configured for an account, the number of segments represents the number of Segments associated with the Account, plus one for the Master view. If no Segments are configured, the segments count will be zero.
- **published hosts** - The number of Published Hosts configured for the Account. A Published Host must be specified for each configuration.
- **origin hosts** - The number of Origin Hosts configured for the Account. An Origin Host must be specified for each configuration.
- **delivery (static) configs** - The number of Content Delivery "static content" configurations created for the Account

- **performance (dynamic) configs** - The number of Orchestrate Performance "websites & apps" configurations created for the Account
- **(chunked) streaming configs** - The number of Chunked Streaming Configurations created for the Account

> **Note:** You can change the selected Company (and Account) using the Company/Account icon at the top right of the page.

Activity and Events

User activity for the current Company (across all Accounts) is shown for the last 12 hours and includes common activities such as Purge request, Configuration changes, and User account updates.

Segments

The most popular segments for the current Company (across all Accounts) are displayed for the current month.

Performance Dashboard

The Limelight Control Performance Dashboard provides a comprehensive snapshot of your company's performance and summarizes information in the Traffic Report and Status Codes Report. The Performance Dashboard is different from the older Dashboard, which is also available.

The Performance Dashboard comprises two sections:

- The top section consists of tabs that show key performance metrics for the Delivery product, such as average bandwidth.

- The lower section provides information about leading metrics, such as locations with the highest data transfer rate. Information is displayed in card format.

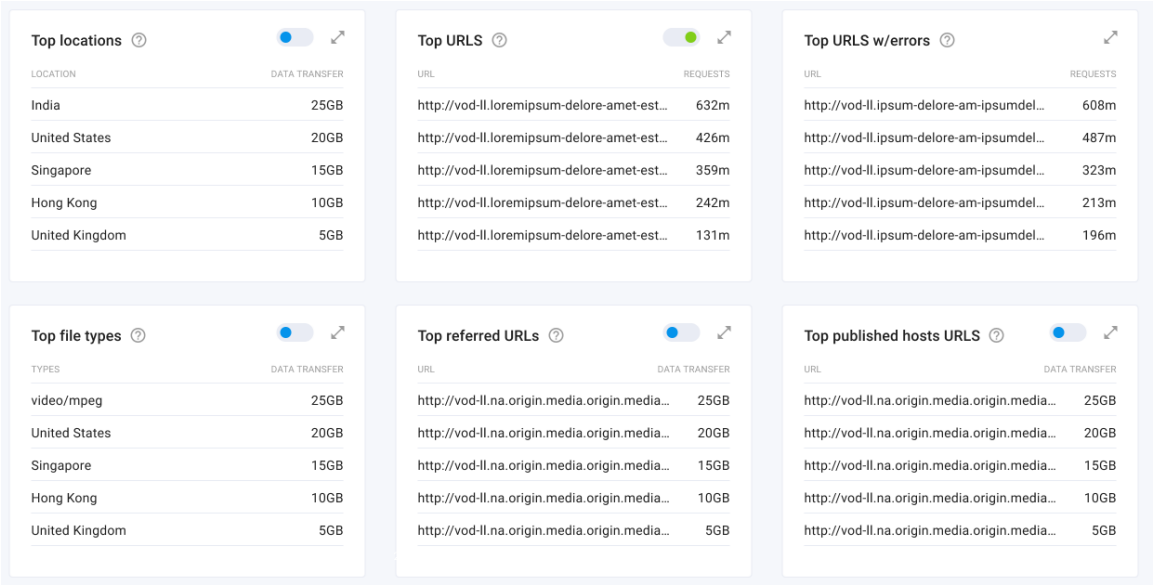

Note: The content of the lower section does not vary when you select a different tab in the top section; the two sections are independent.

By default, data presented is an aggregate of all your company's accounts.

Instructions for using the Performance Dashboard are in these sections:

Displaying the Performance Dashboard

To view the Performance Dashboard, you must set it as your default landing page.

1. After logging into Limelight Control, click the **Profile** icon (see Menus and Links) at the top right of the screen.
2. Select **My Account** from the subsequent drop-down menu.
   The **Edit My Account** page is displayed.
3. Select **Performance Dashboard** from the **Landing Page** drop-down menu.
4. Click the **Save** button.
   When you next log into Limelight Control, the Performance Dashboard is automatically displayed.

Selecting Accounts and Date Ranges

Limit Performance Dashboard data to specific accounts and date ranges by making selections in the drop-down menus at the top of the screen:

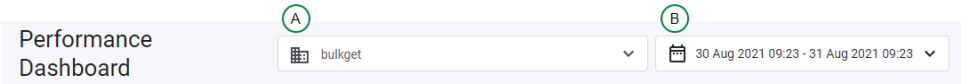

A - Account Selector

B - Date Range Selector

For additional information on selecting date ranges, see Selecting a Date Range.

Key Performance Metrics Tabs

Overview

The following table describes each tab.

Tab	Description	Additional Information
Throughput	Average data transfer rate measured in bits per second	Traffic Report: Overview Tab
Data Transfer	Average amount of data transferred, measured in bytes	Traffic Report: Overview Tab
Requests	Average number of requests	Traffic Report: Overview Tab
Requests Efficiency	How the CDN performed in terms of serving end-user requests from the	Traffic Report: CDN Efficiency Tab

Tab	Description	Additional Information
	cache instead of from the origin. Expressed as a percentage according to the formula: *Responses Served from the Cache/All Incoming Requests + All Outgoing Responses) * 100%*	
Data Transfer Efficiency	How the CDN performed in terms of serving data from the cache instead of from the origin. Expressed as a percentage according to the formula: *Responses Served from the Cache/All Incoming Requests + All Outgoing Responses) * 100%*	Traffic Report: CDN Efficiency Tab
Errors	Summary of requests for content that resulted in the following non-200 status codes: 206 400 403 404 503 Only error codes returned for the selected account and date range are displayed.	Status Codes Report: Overview Tab, URLs Tab
Services	Data transfer rate in bits per second for each service (HTTP and HTTPS) enabled for the account.	Traffic Report: Overview Tab, Details Tab, Geography Tab **Note:** In the Traffic Report, services are referred to as "Protocols."

Note:
If traffic has not been generated for an account, indicators are displayed in the tab headers:

- The words "Not available" are displayed for the **Errors** and **Services** tabs.
- The other tabs display a zero.

Working with Tabs

Each tab displays data for the metric displayed in the tab name.

The average of the tab's metric for the selected accounts and date range is displayed below the tab name. The following example shows that the average Throughput is 192.19 Kbps:

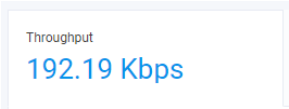

If a metric has not been produced for the selected account and reporting period, the message "Not available" is displayed.

In the **Throughput**, **Data Transfer**, and **Requests** tabs, a toggle on the right above the chart allows you to specify chart content:

Toggle	Description
IN & OUT	Data broken out by **IN** and **OUT** where: • **IN** is traffic coming into the CDN from customer origins. • **OUT** is traffic leaving through published hosts to the requesting client. Chart lines are color coded and identified by labels under the chart.
TOTAL	Sum of **IN** and **OUT**

Configure chart content by making a selection in the toggle.

Leading Metrics Cards

Overview

Working with Cards

Overview

Each Leading Metric card displays the entities with the highest metric for the measure displayed in the card. Each card displays at most five entities; if fewer entities are available, only those are displayed. If more entities are available, you can view them in the card's expanded view (see Working with Cards).

Card	Card Content	Additional Information
Top locations	Geographical locations from which requests for content originated	Traffic Report: Geography Tab
Top URLS*	URLS (non-referrer and non-published) that received requests	Traffic Report: Hosts & URLs Tab

Card	Card Content	Additional Information
Top URLS w/errors*	URLS (non-referrer and non published) that received requests and responded with HTTP error codes	Status Codes Report: Overview Tab, URLs Tab
Top file types*	Sort of file (for example 'text/html', 'text/plan') requested	Traffic Report: Hosts & URLs Tab
Top referred URLs*	Referrer URLs that received requests	Traffic Report: Hosts & URLs Tab
Top published hosts URLs*	Published hostnames that received requests	Traffic Report: Hosts & URLs Tab

*Because this report uses data with daily granularity, date range selections of "Today" and "Last 24 hours" are not available (see Request Proration). Also, only the top 50 URLs per day data is shown. Data for those metrics is only in GMT-7.

Working with Cards

Each card except **Top URLs w/errors** includes a toggle, and all cards include an expand icon:

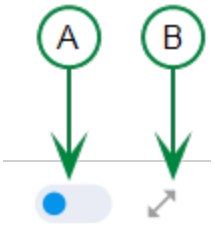

A - Toggle

B - Expand Icon

You can take the following actions:

- Click the toggle to switch between Requests view and Data Transfer view. Text below the toggle changes to reflect the view:

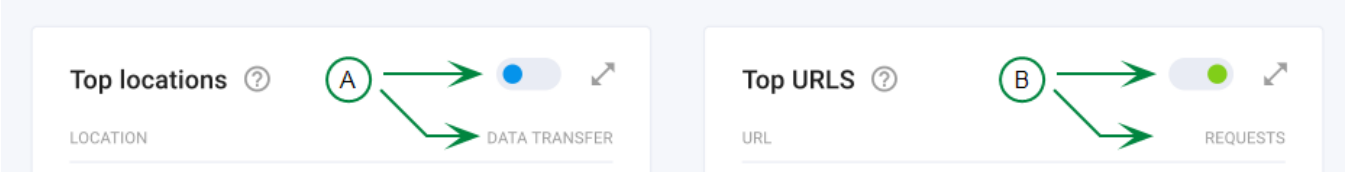

A - Card in Data Transfer view

B - Card in Requests view

- Click the expand icon to switch to expanded view, where data is displayed in a dialog.

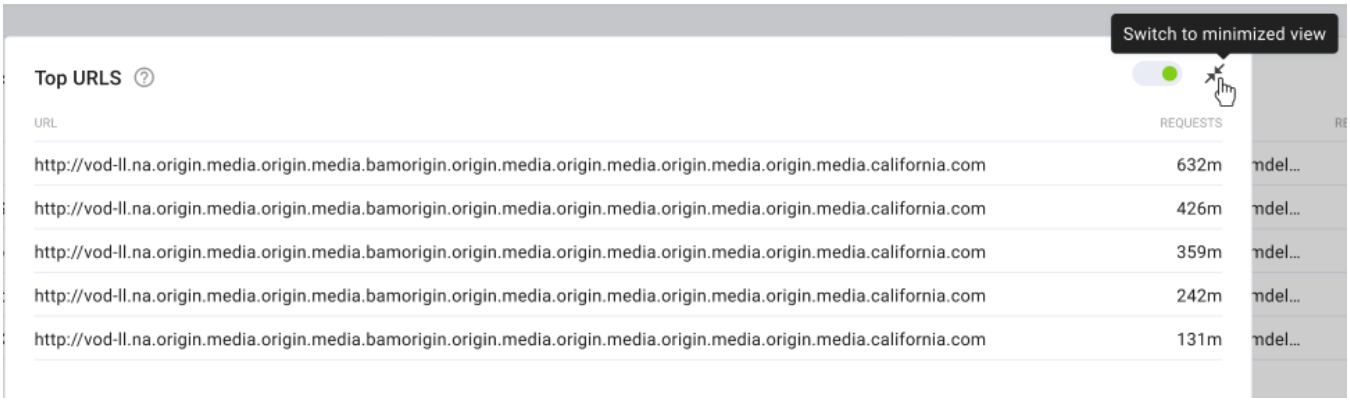

- Click the minimize icon in the dialog to switch back to minimized view.
- Hover the pointer over a URL to view the complete URL.

Note: If a metric has not been produced for the selected account and reporting period, the message "No data available" is displayed.

Configuration

In the *Configure* section of Limelight Control, you can directly create, edit and view the configurations (settings) for many Limelight services, including.

- Content Delivery
  - Delivery
  - Content Security
  - DNS Services
- Video Delivery
  - Live Streaming
  - Chunked Streaming

> **Note:** Some configuration options must be purchased or specifically enabled by Limelight before they become available in self-service configuration. Please contact your Account Manager or Limelight Support if you have questions about access to a specific configuration option.

Configuring *Content Delivery* (read-only)

Content Delivery delivers content via HTTP and HTTPS for all file formats. Both full (entire file) and progressive (range request) downloads are supported.

When you select the **Caching & Delivery** option in the **Configure** menu, the *Configurations* page is displayed.

> **Note:** Caching & Delivery settings are read-only and will be removed in a future release. Please use Caching & Delivery (v2) instead.

The *Configurations* page displays a list of the *Content Delivery* configurations for the currently-selected *Company* and *Account* .

In the *Caching & Delivery* page header, you will see either a *static content* label or a *websites & apps* label, depending on which feature is activated for the currently-selected Account.

The following information is shown for each configuration:

- **Published Host** - The public URL prefix used in links to your published content (URLs seen by end users)
- **Published Path** - The URL path, if any, to use with the **Published Host**
- **Origin Host** - The private URL prefix used by Limelight to retrieve and cache content from your origin server (not visible to end users)
- **Origin Path**- the URL path, if any, to use with the **Origin Host**
- **Host Header** - The value that Limelight will include in the HTTP `Host` header when making requests to your origin
- **Protocol** - The level of HTTP protocol security to use when delivering your cached content to end-users

Creating a New Configuration

To create a new configuration...

- choose whether the configuration is for *static content* or *websites & apps*
- for a new Published Host, click the **new** button at the top of the list
- for a Published Host already in the list, click the **new** button under that host

The *Create configuration* screen will be displayed. After you've filled in the configuration fields in each of the sections, click **Activate** (at the bottom of the page) to enable your new configuration.

Content Location

Setting	Information Requested	Purpose	Selecting the Right Option
Published Protocol	The level of HTTP protocol security to use when delivering your cached content to end-users	To ensure your content is delivered with the level of security you require	• To always deliver content insecurely, select **HTTP** • To always deliver content securely (via SSL), select **HTTPS** • To deliver content using the protocol specified in the incoming HTTP request, select **Both HTTP and HTTPS**
Published Host	The fully-qualified domain name that will be used in all public links (Published URLs) to your cached content **Note:** A URL that includes the Published Hostname is referred to as a Published URL.	To direct your users to the *Content Delivery* service (instead of your origin)	In the **Published Host** field, enter the published hostname specified in the *Welcome Letter* associated with your Limelight Account, or a CNAME if desired. The published hostname provided by Limelight will be in a form similar to: `accountname.vo.llnwd.net` If you prefer to publish under a different hostname, you can use a DNS CNAME record to alias (point) your desired name to Limelight published hostname. **Notes:** • IP addresses are not accepted. You must enter a fully-qualified domain name. • If you can't find the Limelight published hostname in your *Welcome Letter*, please contact Limelight Customer Service If you want to use a directory name "alias" for a particular origin path, you can add the alias by entering it in the **Published Path** field.
Published Path	The "alias" for a particular origin path.		If the path ends with a filename, check the **This path ends with a filename** checkbox. If you want to publish select few types, check the **Only publish files with these extensions** checkbox. Then enter file extensions in the field that displays. When done entering all extensions, press the enter key or tab key.

Location of Content Origin	The location of the content you want the *Content Delivery* service to deliver (the "origin")	The *Content Delivery* service needs to know where to find your content when users first request it, and also when it needs to be refreshed in the cache	If your content is not stored with Limelight, choose **Outside Limelight infrastructure**. Otherwise, choose the appropriate Limelight Storage location. **Note:** If you are using Limelight storage but your storage option is not shown, your *Content Delivery* service is not fully configured. If this is the case, please contact Limelight Customer Service .

Table 1. Configure - Delivery - Content Location Settings

If you choose **Outside Limelight infrastructure** in **Location of Content Origin**, the following additional fields are displayed:

- Origin Protocol
- Origin Host
- Origin Path
- Origin HTTP Port

If you choose a Limelight storage option in **Location of Content Origin**, the following fields are displayed:

- Origin Path

Setting	Information Requested	Purpose	Selecting the Right Option
Origin Protocol	The HTTP protocol(s) to use when retrieving content from your origin (when the content is not found in cache or has expired in cache)	To ensure your content is retrieved with the level of security you require	• To always retrieve content insecurely, select **HTTP** Always • To always retrieve content securely (via SSL), select **HTTPS Always** • To retrieve content using the protocol specified in the user's HTTP request, select **Match Inbound Protocol**
Origin Host	The fully-qualified domain name or IP address of your origin server	The *Content Delivery* service needs to know where to get your content when users first request it, and also when it needs to be refreshed in the cache	Enter the domain name or IP address of your origin server in the **Origin Hostname** field. Please note that if you enter a domain name, it must be fully qualified.
Origin Path			If your content is all in particular path on your origin, or you added a directory name "alias" with the Published Hostname for a particular origin path, you can enter the origin path by clicking the **Add Path** link

Origin HTTP port number	The HTTP port number to use when communicating with your origin server, using the **Origin Host** and **Origin Path** you specified	If you are using a port other than the default (80) for HTTP, the *Content Delivery* service needs to know which port you've chosen	Leave the default port number for HTTP unless you are using another port number. If so, enter the new port number in the **Origin HTTP Port Number** field. **Note:** The default for HTTPS is 443, and this is the value used by Limelight for all HTTP requests to origin (the value is not editable).
Host Header	The value to include in the `HTTP Host` header when communicating with your origin server	To help prevent end users from requesting content directly from your origin.	If you plan to block requests to your origin based on the value of the `Host` header, select **Published Hostname** or enter a value in the **Value** field If you are hosting more than one origin on a single server, please see the additional information below.

Table 2. Configure - Delivery - More Content Location Settings

Host Header Details

Browsers usually include the origin domain name of the requested URL in the HTTP `Host` header. You can use this behavior to detect and block such requests on your origin, denying those with a `Host` header that matches your domain name, and accepting those that match either your **Published Host** or another value you enter in the **Value** field.

If you are hosting more than one origin on a single server and you want to block based on `Host` headers, don't use **Published Hostname** - enter a value in the **Value** field instead. If you are hosting more than one origin on a single server and you don't want to block based on `Host` headers, choose **Origin Host** .

Example Settings

Configuration Field	Value	Notes
Protocol	HTTPS	Accept only HTTPS requests for cached content
Published Host	`published.host.com`	Use a CNAME alias instead of the name provided in the Welcome Letter (need to set up the CNAME separately)
Published Path	`/pubimages/`	Use the pubimages directory to uniquely identify the content in cache
Origin Protocol	HTTP Always	Always use HTTP to communicate with the origin server
Origin Host	`origin.host.com`	
Origin Path	`/images/`	Directory path to the origin content; note that this doesn't need to match the path (if any) for the Published Hostname
Origin HTTP Port	`80`	Use the default HTTP port (no need to change anything)

Configuration Field	Value	Notes
Host Header	Published Hostname	This will block most browser requests made directly to origin

Table 3. Configure - Delivery - Content Location - Example Settings

Using the example configuration settings above, if `favicon.ico` is not cached for this configuration, or has expired in cache, a request to `https://published.host.com/pubimages/favicon.ico` will result in an origin request for `http://origin.host.com/images/favicon.ico`, with an HTTP `Host` header of `published.host.com`.

Caching Rules

Setting	Information Requested	Purpose	Selecting the Right Option
Origin Cache Control & Expiration Header	Whether to override the default method for determining if an object in cache is expired	In some cases you may want to take explicit control over object expiration times (TTL - "Time To Live").	To allow *Content Delivery* to calculate TTL, select **Honor Origin Cache-Control and Expires headers**. Otherwise, choose **Override Origin Cache-Control header and TTL values**. If you want to set TTL to a specific length of time, select one of the times in the Time To Live (TTL) drop-down menu. Otherwise, to allow adaptive TTL calculation, select Custom from the Time To Live (TTL) drop-down menu.
Cache large files on first request	Whether you want any request for an object to force the full object to be cached, even if the request is cancelled.		**Note:** This feature is intended for large file downloads, and is not recommended for caching website objects (such as image, CSS, and JavaScript files).
Ignore "No cache" header	Whether *Content Delivery* should ignore certain `Cache-Control` headers when determining whether or not to cache an object retrieved from your origin	You may want to cache objects regardless of origin settings that attempt to turn caching off	If you want to ignore the following `Cache-Control` headers: • `Cache-control: no-cache` • `Cache-control: no-store` • `Cache-control: private` • `Pragma: no-cache` enable this option. Otherwise, leave it disabled.

Specific Query String Caching	Whether to use URL query terms to determine whether or not objects are cached	You may want to increase cache efficiency by ensuring certain objects are not duplicated due to variations in their query terms	Choose the option that caches the minimum number of objects necessary based on query parameters: • **Strip no query terms from the cache key** • **Strip all query terms from the cache key** • **Exclude specific query terms** • **Keep only specific query terms** **Note:** For the **Exclude specific query terms** and **Keep only specific query terms** options, you must enter a comma separated list of the query terms to be excluded or included
Vary Headers **Note:** for static content configurations only	Which `Vary` response header fields Content Delivery should use when differentiating versions of an object in cache	Content Delivery stores a separate version of a requested object for each unique set of request header fields specified by the `Vary` header. If the `Vary` header specifies request header fields that change frequently, multiple copies of the same object may be stored in cache. To control this behavior, you can configure *Content Delivery* to ignore all `Vary` headers or specific `Vary` headers when caching and retrieving objects. All of the Vary headers associated with the object are still maintained and passed on to the client in the response.	• If you only want to cache a single version of an object regardless of its `Vary` header fields, choose **Ignore all Vary headers** • If want cache a new version of an object whenever any of its `Vary` header fields changes, choose **Do not ignore Vary headers** • If want cache a new version of an object whenever all but certain specified `Vary` header fields change, choose **Ignore specific vary headers** and select the `Vary` headers fields to ignore
Partial Cache	Whether to use Partial Caching to improve cache performance	Partial Caching is a *Content Delivery* feature that caches commonly-requested portions of	To enable this setting, check the **Partial Cache** checkbox, and in the associated field, enter a Regex value that matches the object URLs you

		HTTP GET GET ranges. This optimization can significantly improve performance for large media files.	want to optimize
N Byte Download	Whether to download the first "n" bytes to improve cache performance	*Content Delivery* can automatically cache a specified number if bytes from the beginning of cached files. This optimization can improve first-byte response times in some scenarios.	To enable this setting, check the **N Byte Download** checkbox, and in the associated field, enter a Regex value that matches the object URLs you want to optimize

Table 4. Configure - Delivery - Caching Rules Settings

Use the **Honor Origin Cache-Control and Expires headers** setting unless you have a specific reason to override the way in which object Time To Live (TTL) is calculated.

By default, *Content Delivery* considers an object “stale” (expired from cache) if the number of seconds specified by the associated `Cache-Control: s-maxage` or `Cache-Control: max-age` header has elapsed since initial caching or since the last freshness check, or if neither header is present, if the date and time in the `Expires` header has passed. The order of precedence is `Cache-Control: s-maxage` , `Cache-Control: maxage` , then `Expires`.

If no explicit freshness information is supplied (there are no `Cache-Control: s-maxage,Cache-Control: max-age` or `Expires` headers), and a `Last-Modified` header is present, the CDN will by default use the adaptive cache freshness algorithm to calculate remaining TTL, based on 20% of the age of the cached response, subject to a floor of 3 seconds and a ceiling of 3 days.

If you need to override (ignore) the above behavior, you can use the **Override Origin Cache-Control header and TTL values** option to specify a new TTL value using the **Time to live (TTL)** drop-down menu.

You can also control whether generated responses are cached using **the Cache Generated Responses** checkbox (for the **Custom** option) or **Including Generated Responses** (for other values in the drop-down menu).

Notes:
Generated responses are HTTP responses that are generated dynamically (“dynamic content”). These responses often do not include any of the cache control headers needed to determine TTL, and are not cached by default to avoid caching personalized or user-specific responses.

By default, Limelight defines a generated response as one that is missing all of the following headers:

- `Expires`
- `Last-Modified`
- `Cache-Control: max-age`
- `Cache-Control: s-max-age`

If you choose the **Custom** option for **Time to live (TTL)**, you can change the parameters of the cache freshness algorithm using **Specify custom floor and ceiling cache values**.

If desired, the floor (minimum) can be raised and the ceiling (maximum) can be lowered or raised. If min and max are set equal to each other, the TTL becomes explicit, rather than adaptive.

Arc Light

You can use Arc Light to customize how Content Delivery reacts to HTTP requests and responses. Rules can be triggered when a request or response meets pre-defined conditions, such as a pattern match. Rules are designed based on specific customer needs.

> **Note:** This option is available only for *websites & apps* configurations

Configuration Overview

Use Arc Light to customize how Content Delivery reacts to the following HTTP request and response types:

- Requests
  - Any
  - Origin only
  - Edge only
- Responses
  - Any
  - Origin only
  - Client only

For each of the above request and response types, you can assign one rule. Content Delivery will then execute that rule each time it receives the associated request or response type. The rule will be executed on the Edge Server that receives the request or response.

To enable Arc Light for a specific request or response type, check the checkbox next to the desired type (example: ☒ **Rules on Edge Request**). To assign a rule, click one of the rules in the list below the request/response type.

Rules can be triggered when a request or response meets pre-defined conditions, such as a pattern match with:

- The URL, file name, or query term
- The IP address
- The value of a specified HTTP header
- A cookie
- The geographic location of a request (using the IP address)

When a rule is triggered, it can perform a variety of actions, such as:

- Controlling which CORS headers are sent in response to a client request
- Adding a cookie that contains geolocation information
- Adding specific HTTP headers
- Appending special “keys” to cache keys Enabling or disabling GZIP compression
- Controlling whether the requested content is cached and setting content TTLs

Rules are designed based on specific customer needs. If you need to use Arc Light or want more information on the types of rules that can be created, please contact your Account Manager or Solutions Engineer.

Configuration Settings

You can configure these settings:

Setting	Information Requested	Purpose	Selecting the Right Option
Which rules do you want to enable?	If you want to create a new rule, the type of HTTP request or response to associate it with	Content Delivery can trigger rules for several types of requests and responses	(see the options below)
Rules on Any Request	Request type	Content Delivery can trigger rules for several types of requests	To trigger a rule on any type of request received by a Limelight Edge Server, check the **Rules on Any Request** checkbox, and select one of the predefined rules in the list
Rules on Edge Request	Request type	Content Delivery can trigger rules for several types of requests	To trigger a rule on client requests to a Limelight Edge Server, check the **Rules on Edge Request** checkbox, and select one of the predefined rules in the list
Rules on Origin Request	Request type	Content Delivery can trigger rules for several types of requests	To trigger a rule on Limelight requests to your Origin, check the **Rules on Origin Request** checkbox, and select one of the predefined rules in the list
Rules on Any Response	Response type	Content Delivery can trigger rules for several types of responses	To trigger a rule on any type of response received by a Limelight Edge Server, check the **Rules on Any Response** checkbox, and select one of the predefined rules in the list
Rules on Origin Response	Response type	Content Delivery can trigger rules for several types of responses	To trigger a rule on responses received from your Origin, check the **Rules on Origin Response** checkbox, and select one of the predefined rules in the list
Rules on Client Response	Response type	Content Delivery can trigger rules for several types of responses	To trigger a rule on responses received from the requesting client, check the **Rules on Client Response** checkbox, and select one of the predefined rules in the list

Media Delivery

Content Delivery supports "seeking" or "scrubbing" (skipping back and forth) within FLV and MP4/H.264 video files. Seeking is controlled via parameters specified in the query terms of the request URL.

Setting	Information Requested	Purpose	Selecting the Right Option

Enable FLV Scrubbing	Whether to allow video client to skip forward and back (seek) within FLV files based on parameters specified in the query terms of the request URL.	Custom clients may want to provide the “seek” capability (“forward” and “back” buttons)	To enable this feature, check the **Enable FLV Scrubbing** checkbox
Enable MP4/H.264 Scrubbing	Whether to allow video client to skip forward and back (seek) within properly segmented MP4 files based on parameters specified in the query terms of the request URL. When you use this option, any query terms in the URL are ignored. **Note:** Query terms are interpreted by EdgePrism, and influence what part of an MP4 file is presented to a user. Other query terms in the URL may be ignored, which may influence the behavior of the origin that delivers the file.	Custom clients may want to provide the “seek” capability (“forward” and “back” buttons)	To enable this feature, check the **Enable MP4/H.264 Scrubbing** checkbox

Table 5. Configure - Delivery - Media Delivery Settings

Optimization

Setting	Information Requested	Purpose	Selecting the Right Option
Type of Compression	Whether to use Gzip compression when delivering XHTML, JavaScript, CSS, and other text files	Compressed objects are delivered more quickly, potentially improving the user experience	• If you want to provide all compressed files from your origin server, choose the **Gzip Passthrough** option • If you prefer to have the *Content Delivery* service compress files when the requesting client can accept them, choose **Gzip on-the-fly** • If you need to modify Gzip compression defaults, choose **Custom**, then either **Gzip on-the-fly** or **Gzip Passthrough** , and enter your Gzip modification extensions • You can also choose **No compression** if none of your files should be delivered compressed • For more information on this feature, see Gzip Details
TCP Acceleration	The “profile” to use when accel-	In certain circumstances, you may want to change the	When **TCP Acceleration** is enabled, the XDLL profile is the

	erating the transfer of IP packets by modifying default TCP parameters	**TCP Acceleration** profile to optimize your delivery performance	most efficient in many cases. **Note: TCP Acceleration** is an advanced configuration setting, and should only be changed if you're an expert user.
Enable chunked response to client **Note:** for static content configurations only	Whether the *Content Delivery* service can maintain open TCP connections to your origin server	If there is a cache "miss" and your origin doesn't provide a `Content-Length` header, this option allows *Content Delivery* to serve the requested content more efficiently (in "chunks")	We recommend you enable this option. Otherwise, new TCP sessions must be established for each new request to origin, and cache miss requests are delivered only when the entire object has been transferred from origin.

Table 6. Configure - Delivery - Optimization Settings

Gzip Details

When **Gzip Passthrough** is enabled, and a client indicates (via HTTP request header) that it prefers to receive compressed content, the *Content Delivery* service will serve a compressed version of the requested object if one is available on the origin server.

Note: ***Gzip Passthrough*** *is available to all customers. If it is not enabled for you, please contact Limelight Support.*

If **Gzip On-the-fly** is selected, the *Content Delivery* service creates, caches, and delivers Gzip-compressed content as needed.

Compressible file types include: `action, ashx, asmx, asp, aspx, axd, cfm, css, css3, csv, do, doc, docx, htm, html, js, jsf, json, jsp, php, portal, rtf, svg, svgz, tsv, txt, xhtml, xml, site root (/),` and extensionless URLs.

Request & Response Headers

Setting	Information Requested	Purpose	Selecting the Right Option
Client Analytics	Whether you want the *Content Delivery* service to provide geographic user information when requesting content from your origin	You may want to internally capture, analyze and report on user geographic information.	To use this feature, check the **Client Analytics** checkbox. The geo information is provided to your origin server via two request headers: `X-IP-Geo-Country` and `X-IP-Geo-All`. The geo fields provided are `continent` , `country` , `state` , `city` , `dma_id` , and `asn`.

Add client IP address to origin request header **Note:** for static content configurations only	Whether you want the *Content Delivery* service to provide the requesting client's IP address in a custom header when requesting content from your origin	You may want to internally capture, analyze and report on user IP information	To enable this feature, check the **Add client IP address to origin request header** checkbox, and enter the header name(s) that should contain the client IP address. The default header name is `True-Client-IP`. Note that the above headers are in addition to `X-Forwarded-For`, which is always provided to the origin.
POST Requests	Whether you want to accept or ignore POST requests from clients	If you are using a custom client to display content, you may want to allow it to communicate analytics or other information to your origin. Alternatively, you may want to convert POST requests to GET requests, or simply ignore them.	• To ignore all POST requests, select **Disable HTTP POST requests**. *Content Delivery* will respond with an HTTP `413 Request Entity Too Large` status code to all POST requests. • To accept POST requests and pass them through to your origin, select **Enable HTTP POST requests**. If a POST request body exceeds 500 MB, *Content Delivery* will respond with an `413 Request Entity Too Large` status code. • To accept POST requests but treat them as GET requests, select **Enable HTTP POST requests,** and check the **Discard request body on POST request** checkbox. POST bodies will be discarded.
Add custom request header	Whether you want to include custom headers and values whenever *Content Delivery* makes a request to your origin	If you want to tag all requests from *Content Delivery* for later analysis	To add a custom origin request header, click the **Add custom request header** link, and enter a unique header name and value
Add Limelight server IP address when respond-	Whether to	If you are using a custom	To enable this feature, check the

ing to client	provide clients with the IP address of the *Content Delivery*Edge Server responding to their requests	client to display content, and you are also capturing performance-related data via the client, you may want to include the *Content Delivery*Edge Server IP address for later analysis and reporting. The IP address will be provided in the X-IP-Address response header.	**Add Limelight server IP address when responding to client**checkbox
Add custom response header	Whether you want to include custom headers and values whenever *Content Delivery* responds to a client request	If you are using a custom client to display content, you may want to provide it with information that uniquely identifies the *Content Delivery* service, Limelight Account, etc.	To add a custom client response header, enter a unique header name and value and. Click the "+" button to add additional headers.
Enable Custom Debug Headers	Whether you want to enable Custom Debug Headers	By making an HTTP content request with special "Custom Debug Headers," including a shared secret specific to your service, you can retrieve cache-related information about individual content objects and prevent others from accessing the information.	In the **Debug Headers** field, enter one or more "tags" to include in the Custom Debug Headers. Then in the **Secret Key To Request Debug Information** field, enter the secret key (shared secret) provided by Limelight when the Custom Debug Headers feature was enabled. For more information, see Secure Cache Diagnostics.

Table 7. Configure - Delivery - Request & Response Headers Settings

Secure Cache Diagnostics

When troubleshooting caching issues, customers can now directly access diagnostic information about cached content.

Content Delivery customers with *Configuration Self Service* can now request cache-related information about individual content objects, without the risk of this information being accessible by others.

To enable this feature, check the *Enable Custom Debug Headers* checkbox in the *Request and Response Headers* section of *Content DeliveryConfiguration Self Service*, and provide a comma-separated list of object properties that should be returned in the response, along with a Secret Key to authenticate the request.

Diagnostic response headers can include the following information:

- Whether or not a response is cacheable
- How the cache responded to a request (hit, miss, etc.)
- The number of seconds before the cached response will be considered stale (TTL)
- The total number of seconds representing the freshness lifetime of the response (age + TTL) and how the value was determined (headers, overrides, adaptive TTL, etc .)

When the feature is activated, you will be provided with a unique shared secret.

The properties that can be requested, and their associated response headers and values, are:

Request Key	Response Header	Return Values

`is-cacheable`	`X-LLNW-Dbg-Is-Cacheable`	`Yes`
		`No`
		`Negative`

`cache-hit-type`	`X-LLNW-Dbg-Cache-Hit-Type`	`HIT`
		`MISS`
		`REFRESH_HIT`
		`REF_FAIL_HIT`
		`REFRESH_MISS`
		`CLIENT_REFRESH_MISS`
		`IMS_HIT`
		`NEGATIVE_HIT`
		`DENIED`
		`OFFLINE_HIT`
		`REDIRECT`

`ttl`	`X-LLNW-Dbg-TTL`	`n{...} seconds` an integer followed by a space and the string "seconds"
`fresh-life-total`	`X-LLNW-Dbg-Fresh-Life-Total`	`n{...} seconds` an integer followed by a space and the string "seconds"

If the secret is invalid, the `X-LLNW-Dbg-Hdrs` header will be ignored and the request will be processed without it.

Request & response example:

Request	Response
`GET http://www.customer.com/object.txt HTTP/1.1...X-LLNW-Dbg-Hdrs: is-cacheable,cache-hit-typeX-LLNW-Dbg-Secret: sharedsecret`	`HTTP/1.1 200 OK...X-LLNW-Dbg-Is-Cacheable: Yes...X-LLNW-Dbg-Cache-Hit-Type: HIT`

Failover

Normally, when the CDN receives an HTTP `404 (Not Found)`, `503 (Service Unavailable)` or `504 (Gateway Timeout)` response from your origin, the error is passed back to the requesting client. You can modify this configuration option as follows:

- For `404` errors:
  - Serve "stale" content from the CDN Cache, or
  - Request content from a "backup host" (with or without a path), or
  - Redirect to a custom "Not Found" URL.
- For `503` and `504` errors:
  - Request content from a "backup host", or
  - Redirect to a custom "Service Unavailable URL"

Notes:
- Failover URLs must match their own configuration within the CDN
- For `404` error redirects, the original request is reissued to the fallback URL with any modifications still in place
- `503` or `504` errors may have been generated by the origin, but could also be generated by CDN if a connection can't be made to your origin

Setting	Information Requested	Purpose	Selecting the Right Option
Serve stale content instead of 404 error	If the requested content is cached but stale (expired), and there is an HTTP `404` status when requesting a fresh version from your origin, whether you want to pass the `404` status back to the client, or serve the stale content instead	If an object has expired in cache, and your origin server returns a `404 (Page Not Found)` error when *Content Delivery* attempts to get a fresh copy of the object, you may want to serve the expired object instead of allowing the client to handle the `404` message.	If it's not acceptable for the client to handle the `404` message, and you are OK with serving stale content instead, check the **Serve stale content instead of 404 error** checkbox. Note that if there is no cached object, a 404 message will still be returned to the browser.
Request content from backup host on 404 error	If there is an HTTP `404` status when requesting fresh content from your origin, whether to try a backup origin (hostname only) before handling the `404` status	If your primary origin returns a `404` status, and you have a backup origin, you may want *Content Delivery* to try the backup before handling the error	To try a backup origin if the primary origin responds with a `404` status, enter the fully-qualified hostname of the backup origin. **Note:** Specific ports are not supported.
Use custom "Not Found" page	Whether you want to pass HTTP `404` status messages back to the client, or serve a custom error page instead	If an object has expired in cache, and your origin server returns a `404` error to *Content Delivery* , you may want to serve a custom error page instead of allowing the client to handle the `404` message.	If you want to take control over the content displayed by clients when there is a `404` from origin, enter the fully-qualified URL of the content to serve.

Request content from backup origin URL on 404 error	If there is an HTTP `404` status when requesting fresh content from your origin, whether to try a backup URL path before handling the `404` status	status, and you have a backup origin, you may want *Content Delivery* to try the backup before handling the error	To try a backup URL path if the primary origin responds with a `404` status, enter the fully-qualified path on the backup origin. **Notes:** • You can specify either the HTTP or HTTPS protocol, and a port number if desired. • This option is required when using the Intelligent Ingest feature of Origin Storage
Request content from backup host on 5xx error	If there is an HTTP `5xx` status when requesting fresh content from your origin, whether to try a backup origin before handling the `5xx` status	If your primary origin returns a `5xx` status, and you have a backup origin, you may want *Content Delivery* to try the backup before handling the error	To try a backup origin if the primary origin responds with a `5xx` status, enter the fully-qualified hostname of the backup origin. **Note:** Specific ports are not supported.
Use custom "Service Unavailable" page	Whether you want to pass HTTP `503` and `504` status messages back to the client, or serve a custom error page instead	If an object has expired in cache, and your origin server returns a `503 Service Unavailable` or `504 Gateway Timeout` error to *Content Delivery* , you may want to serve a custom error page instead of allowing the client to handle the error message.	If you want to take control over the content displayed by clients when there is a `503` or `504` error from origin, enter the fully-qualified URL of the content to serve.

Table 8. Configure - Delivery - Failover Settings

Content Security

IP Access Control

Setting	Information Requested	Purpose	Selecting the Right Option
Enable IP Access Control	Whether you want to "allow list" or "deny list" requests based on IP address lists and IP-based geographic locations	IP Access Control allows you to exclude specific geographies or limit access to known entities	Assign access lists to the Caching & Delivery configuration using the following drop-down menus: • **By IP address list**: Select one or more existing lists, then choose either **Deny** or **Allow** to indicate the type of restriction. Click **Add** to add the lists to the **Access control list for this configuration** section. • **By geolocation**: Select one or more geographic areas (continents or countries), then choose either **Deny** or **Allow** to indicate the type of restriction. Click **Add** to add the lists to the "Access control list for this configuration" section. **Access control list for this configuration** Section You can select a default security setting for the configuration - either **Default Allow** or **Default Deny**. You can then add one or more IP address lists and geographic locations that modify the default setting. IP address lists and geolocations can be "mixed and matched" in any order desired. To move an item in the list, move the mouse pointer over the item and use the vertical ellipses to drag and drop the item to another location in the list. If you have the correct permissions, click **Manage IP Lists** to display a dialog that allows you to create new IP access lists. You can also view, edit, and delete existing IP address lists. To view list details, click the **+** icon to the left of a list. • The text "Used by configs in accounts" shows which Accounts have configurations that use the list. • The text "Limited to accounts" shows any accounts to which your *Company Admin* has limited the list. To create a new list, click the **new list** button at the top of the dialog, then: 1. Provide a name for the list. 2. Provide a single IP address or range of IP addresses. You can also create and upload CSV files of IP addresses. Click the link to see a sample CSV file. 3. Optionally limit the list to accounts. 4. Click the **Save** button. The new list is now available in the **By IP address list**: drop-down menu at the top of the section. To deny access to end users attempting to access your content using an anonymous VPN from an unauthorized geolocation. select the **Deny 'Anonymized with VPN' access** option. **Notes:** • IP address lists and geographic locations are processed in the order they are specified (top to bottom). Once a match is found, subsequent lists and locations are ignored • Users with the *Company Admin* role can manage lists for all accounts. Users with the *User* role who have been granted the *Manage Delivery Configurations* permission can apply all lists in the Accounts for which they have been

Setting	Information Requested	Purpose	Selecting the Right Option
			granted management permission. • Changes made to IP address lists are applied immediately and affect all Account configurations which use them (even legacy configurations that can't be edited in Control. • IP address lists cannot be deleted if they are in use.

Table 9. Configure - Delivery - Content Security Settings

Content Delivery makes it easy to implement and manage IP-based access control using both IP addresses and geographic locations.

Content Delivery Configuration Self Service provides access control using IP addresses and geographic locations ("geo-fencing"). When configuring an Account, you can associate lists of IP addresses and groups of geographic locations with the Account and specify whether to allow or deny each. When managing IP address lists, you can also view whether they are currently in use and which Accounts they are associated with (or limited to).

MediaVault

Setting	Information Requested	Purpose	Selecting the Right Option
Enable MediaVault content protection	Whether you want to use *MediaVault* to provide additional content security. *MediaVault* provides high-performance URL authentication.	*MediaVault* can help you prevent “deep linking” and other unauthorized viewing behavior	To enable this feature, check the **Enable *MediaVault* content protection** checkbox, and provide a primary and secondary “shared secret” (both used to prevent URL tampering). You can also change the HTTP Error Code returned by MediaVault from the default 400 code by entering a new value in the **Deny Status Code** field. For more information, see *MediaVault* Details.

MediaVault is a high-performance URL authentication service. *MediaVault* ’s main purpose is to help you secure your content from unauthorized viewing.

MediaVault maximizes authentication performance by using tokens to avoid three-way handshakes (common to other methods of authentication) that can lead to severe connection time latency.

Please note that *MediaVault* is *not* a replacement for DRM and should not be associated with user authentication.

MediaVault works like this:

- You enter a shared secret during the configuration process
- You then generate a token (MD5 hash) for each published URL, based on the shared secret, and append it to the URL in a query term or provide it in a cookie. You can generate the token manually by navigating to the *Configure* > *MediaVault* in the navigation pane, or by creating server-side code on your origin.
- When a request is received, *MediaVault* uses the same hash algorithm to create it’s own token, which should be identical to the one you appended.

- If the tokens match, *MediaVault* then looks for additional *MediaVault* -specific query terms (such as end date/time and IP address/mask) to determine whether the request is valid. If the tokens don't match, the URL was tampered with and the request is rejected.

For more information, see the MediaVault User Guide by navigating to Help Center > Documentation > Delivery > Guides > MediaVault in the navigation pane.

> **Note:** *MediaVault* does not support IPv6-based Client IP or IP Range restrictions at this time.

Advanced

You can use the **Additional Options** step to view any advanced *Content Delivery* configuration changes Limelight makes to your configuration.

If one or more such configurations is changed from its default value by Limelight, the **Additional Options** tab becomes visible, and the advanced configurations and their settings are displayed:

Setting	Information Requested	Purpose	Selecting the Right Option
(various)	(none) This is a read-only display of advanced *Content Delivery* configuration changes Limelight has made to your configuration	The information in the Additional Options step can help you better understand your configuration.	If you have questions about any settings in Additional Options, please contact your Account Manager or Limelight support.

Table 10. Configure - Delivery - Advanced Settings (Read Only)

The advanced configuration options which can be configured for you by Limelight (and become visible in the *Additional Options* step) include:

Option Name	Description
Assume cacheable pending origin response	If an origin request is pending for an object, continue serving the object from cache
Cache entire object if range request less than offset	Cache the entire object for Range requests ending before the specified Byte offset
Cache hit/miss response trigger	Returns HIT or MISS in the X-CDN-Cache response header when the specified request header (trigger) is present
Cache only "popular" objects	Cache only objects that are "popular" based on the specified "points" (the approximate frequency an object is requested, in seconds)
Convert URL ranges to Range requests	Convert URLs ending in /range/x-y or /range/x- to origin GET range requests
Deny requests with specified Referer header(s)	Deny requests with the specified Referer header(s)
Disable object caching	Do not cache objects
Disable persistent	Disable persistent origin connections ("enabled" is the default global configuration)

origin connections	
Do not add max-age on all requests to origin	Don't add Cache-Control: max-age=259200 header on origin requests (but do include any existing Cache-Control headers)
Gzip compression level	Set the Gzip compression level (0 to 9). The default (and recommended) level is 1.
Ignore bad status codes from origin	Ignore bad status codes from origin (40x and 5xx). If FALSE, other rewrite options may redirect the client to specific URLs based on the status code.
Lowest allowed rate-limiting bitrate	Set the lowest bitrate allowed when rate limiting, in KBytes/second
Make cached URLs case-insensitive	Make the URLs of cached objects case-insensitive by converting all characters to lowercase in the Cache Key. When using this feature, all Purge requests must not contain any uppercase characters.
Max duration client can be idle while receiving response	After this time passes, the client is disconnected and the request is aborted. The default is 30 minutes.
Maximum object TTL	Set the maximum TTL value for all cached objects, in seconds, but honor Cache-Control headers if present
Minimum object TTL	Set the minimum TTL value for all cached objects, in seconds, but honor Cache-Control headers if present
Object TTL for "negative" origin response	Set the object TTL, in seconds, when there is a negative origin response (status codes other than 200, 203, 300, 301 and 401 and/or Cache-Control or Pragma headers with certain values). This rewrite overrides other origin cache control headers.
Origin connect timeout duration	Set the timeout, in seconds, for initiating origin connections (how long to wait when trying to establish a connection)
Origin reply timeout duration	Set the timeout, in seconds, for origin replies (how long to wait for a reply from origin)
Persistent client connection duration	Set the duration, in seconds, of persistent client connections
Persistent origin connection duration	Set the duration, in seconds, of persistent origin connections
Redirect clients to source URL	Redirect clients to the source URL with the specified status code
Refresh-check cached content on every request	Check for fresh origin objects (newer versions of objects) on every request. Most commonly used in conjunction with *Ignore bad status codes from origin* to enable the origin to allow or deny every request by inspecting all request parameters, including Cookies.
Remove specified response header(s)	Remove origin response headers that match the specified value
Retry failed *MediaVault* HTTPS hash checks	If an HTTPS *MediaVault* hash check fails, retry the same hash-check URL using HTTP
Store *MediaVault* hash in cookie	Keep the *MediaVault* hash secret in a browser cookie (rather than in a URL parameter)
Treat empty responses with 200 status as 404 status	Treat "empty" origin responses (no content body) with 200 status codes as if they are 404 status codes
Do not apply MediaVault on URLs	Configure MediaVault to ignore URLs matched by the displayed regular expression

matching the regex	

Table 11. Configure - Delivery - Wizard - Additional Options Available

Logging

Setting	Information Requested	Purpose	Selecting the Right Option
Log cookies	Whether you want *Content Delivery* to *stop* saving cookie information in your log files	If you process log files and don't need the information in the Cookie header, you may want to remove it to simplify processing and/or reduce log file size.	If you know you need Cookie header information in your log files, check the **Log cookies** checkbox. Otherwise, leave it unchecked. When this setting is enabled, *Content Delivery* logs all `Cookie` header information, up to a maximum of 8 KB for the entire header (regardless of the number of cookies in the header).
Log request header	Whether you want *Content Delivery* to *start* saving specific Request Headers in your log files	If you process log files and need access to information in the Request Headers, you may want to enable this option	If you know you need Request Header information in your log files, check the **Log Request Header** checkbox and enter the names of the specific headers to log. Otherwise, leave it unchecked.

Table 12. Configure - Delivery - Logging Settings

Notes

You can use the **Notes** field to additional information for others (why the configuration changes were made, etc.). Users can refer to the notes later when browsing historical configuration changes

Previewing a Configuration

 To preview the settings associated with a configuration, click the "eye" icon at the bottom right of the configuration row. For information on the individual settings displayed, please see the setting descriptions in Creating a New Configuration.

Editing a Configuration

 To edit a configuration, click the "pencil" icon at the bottom right of the configuration row. For information on the individual settings displayed, please see the setting descriptions in Creating a New Configuration.

> **Note:** For some configurations created for you by Limelight, only the preview ("eye") icon will be visible. If this is the case, and you need to make changes, please contact your Account Manager or Solutions Engineer. Limelight can continue to manage the configuration, or it can be made available for you to edit in the *Limelight Control*.

Cloning a Configuration

 To clone (make a copy of) a configuration, click the "copy" icon at the bottom right of the configuration row. When you have finished making changes to the settings, click **Activate** to enable the new configuration.

> **Note:** For some configurations created for you by Limelight, only the preview ("eye") icon will be visible. If this is the case, and you need to make changes, please contact your Account Manager or Solutions Engineer. Limelight can continue to manage the configuration, or it can be made available for you to edit in the *Limelight Control*.

Deleting a Configuration

 To delete a configuration, click the "trash" icon at the bottom right of the configuration row.

> **Note:** For some configurations created for you by Limelight, only the preview ("eye") icon will be visible. If this is the case, and you need to make changes, please contact your Account Manager or Solutions Engineer. Limelight can continue to manage the configuration, or it can be made available for you to edit in the *Limelight Control*.

Reverting to a Previous Configuration

 Each time you update a configuration, a new version is assigned.

To revert to a previous configuration:

1. Click the "undo" icon at the bottom right of the configuration row.
   A list of previous versions display in a dialog.
2. Select the version to which you want to revert

   > **Note:** Although you intend to revert to a previous version, the reverted version will become the current version, which will have a new version number. The new version number is displayed at the bottom of the dialog.

3. Click the **Activate** button.

Configuring *Content Delivery* (v2)

Content Delivery delivers content via HTTP and HTTPS for all file formats. Both full (entire file) and progressive (range request) downloads are supported.

Navigate to **Configure** > **Caching & Delivery (v2)** in the navigation pane. The *Caching & Delivery (v2) for* page is displayed.

Configuration List

The *Caching & Delivery (v2) for* page displays a list of the *Content Delivery* configurations for the currently-selected *Company* and *Account*.

The following information is shown for each configuration:

- **Source Hostname** – The private URL prefix used by Limelight to retrieve and cache content from your origin server (not visible to end-users)
- **Source Path**- the URL path, if any, to use with the **Source Hostname**
- **Protocol** - The level of HTTP protocol security to use when delivering your cached content to end-users
- **Published Path** - The URL path, if any, to use with the **Published Host**

Filtering the List of Configurations

Use the **Filter by** drop-down menu and the filter text field to filter the list by specific fields:

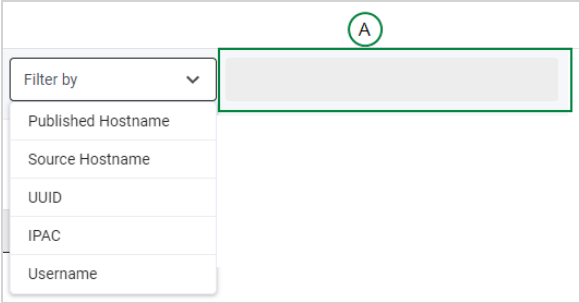

A - filter text field, initially disabled

To filter the list:

1. Make a selection in the drop-down menu.
2. Enter a value in the filter text field.
3. Press the Enter key on your keyboard.
   The list is reduced to include only configurations that match the filter.

Display the original list by clicking the **x** icon in the filter text field:

Read-Only and Hidden Capabilities

For particular use cases, configurations may have certain fields presented as either read-only or hidden (masked).

An entire configuration may be read-only. Or, an editable configuration may have source and/or published fields hidden or read-only.

Source and published fields are:

- Source hostname
- Source path
- Published hostname
- Published path

Hidden fields are masked with asterisks.

Use Case	User Capabilities
Entire configuration is read-only.	Users can only view the configuration. They cannot edit, clone, or revert the configuration.
Configuration is editable, but specific fields are read-only and/or hidden.	Users can modify all but read-only/hidden fields. Users can edit, clone, and revert the configuration.

Examples

Masked Field

Configuration Is Editable and Has Two Read-Only Fields

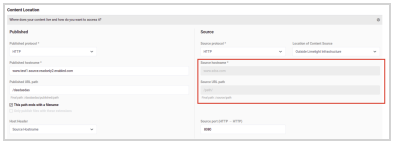

Creating a New Configuration

To create a new configuration, click the **+new** button, and the *Create configuration* screen will be displayed.

Service Profiles

Each new configuration is based on a Service Profile. Service Profiles define the configuration structure and specify default and mandatory options that must be applied on every configuration. A Service Profile can serve as both a guide and a guardrail for the type of content your configuration will serve (characterized by a Use Case).

The **Use Case** and **Service Profile** drop-down menus are disabled:

- In existing configurations.
- After you have selected a Published and Source Protocol while you are creating a new configuration.

> **Note:** If you have not already saved the new configuration but you want to choose another Service Profile, you can do so by exiting out of the **Create configuration** screen and creating a new configuration by clicking the **+new** button.

If you wish to modify a Service Profile or migrate, add, or remove a Protocol Set for an existing configuration, contact your Account Manager.

Page Organization

Configuration options are grouped into sections by functional category, such as *Content Location*. Each section displays all of the related options made available by the Service Profile. For sections other than *Content Location*, the most commonly used options are always displayed. Any remaining options are grouped under the **Advanced**drop-down menu in each section.

Initially, only the Content Location section is visible. Once you select both a **Published protocol** and a **Source protocol**, the rest of the sections and configuration options become available. The combination of a **Published protocol** and **Source protocol** is known as a "protocol set."

The following sections describe the fields on the page:

Content Location	Caching Rules	Arc Light	Media Delivery	Optimization
Headers & Methods	Secure Cache Diagnostics	Failover	Content Security	Logging
Cookie Handling	Redirect	Others	Additional Options	Notes

After you've filled in the configuration fields in each section, click **Activate** (at the bottom of the page) to enable your new configuration.

Content Location

Setting	Information Requested	Purpose	Selecting the Right Option
Published protocol	The level of HTTP protocol security to use when delivering your cached content to end-users	To ensure your content is delivered with the level of security you require	The **Published protocol** and **Source protocol** drop-down menus are disabled: • In existing configurations. • After you have selected a Published and Source protocol while you are creating a new configuration. **Note:** If you have not already saved the new configuration but you want to choose another Published and Source protocol, you can do so by exiting out of the *Create configuration* screen and creating a new configuration by clicking the **+new** button.
Published hostname	The fully qualified domain name that will be used in all public links (Published URLs) to your cached content **Note:** A URL that includes the Published Hostname is referred to as a Published URL.	To direct your users to the *Content Delivery* service (instead of your origin)	In the **Published hostname** field, enter the published hostname specified in the *Welcome Letter* associated with your Limelight Account or a CNAME if desired. The published hostname provided by Limelight will be in a form similar to: `accountname.vo.llnwd.net` If you prefer to publish under a different hostname, you can use a DNS CNAME record to alias (point) your desired name to Limelight published hostname. **Notes:**

			• IP addresses are not accepted. You must enter a fully qualified domain name. • If you can't find the Limelightpublished hostname in your *Welcome Letter*, please contact Limelight Customer Service. If you want to use a directory name "alias" for a particular origin path, you can add the alias by entering it in the **Published URL path** field. If needed, you can add a regex expression to the start of the **Published hostname** field, but you must have permissions to do so. Without the permissions, you are restricted as follows: • When creating or cloning a configuration, you cannot add regex to the field. • If a configuration has regex in the field, you cannot clone the configuration. • When updating a configuration that has regex in the field, you cannot modify any part of the Published hostname. Please contact your account manager if you need assistance with any of these operations.
Published URL path	The path portion of a published URL	To allow your published hostname URL to be more specific and include a path.	Enter the path enclosed in forward slashes. **Notes:** • This field must contain a value and defaults to / • If you have the PERMISSION_CONFIGURE_SSUI_REGEX permission, you can include regular expressions in the path.
This path ends with a filename	Whether the last component in the path is a file	File names are not validated by extension, so when the Published URL path or Source URL path does not end with a slash, it is treated as a file name. **Note:** This field and **Only publish files with these extensions** are mutually exclusive.	Place a checkmark in the checkbox if the path ends in a file name. **Note:** If you check this option, you must make entries in the **Published URL path** and **Source URL path** fields
Only publish files with these extensions	File extensions to publish	Provides flexibility, allowing you to specify file extensions to publish. More flexible than using **This path ends with a filename**, which allows you to specify only one file. **Note:** This field and **This path ends with a**	Place a checkmark in the checkbox, then enter file extensions (excluding a leading period) in the field below the checkbox.

		filename are mutually exclusive.	
Host Header	The value to include in the `HTTP Host` header when communicating with your origin server	To help prevent end-users from requesting content directly from your origin.	If you plan to block requests to your origin based on the value of the `Host` header, select **Published Hostname** or enter a value in the **Value** field If you are hosting more than one origin on a single server, please see the additional information below. For more information, see Host Header Details.
Source protocol	The HTTP protocol(s) to use when retrieving content from your origin (when the content is not found in the cache or has expired in cache)	To ensure your content is retrieved with the level of security you require	See Published Protocol.
Source hostname	The fully qualified domain name or IP address of your origin server	The *Content Delivery* service needs to know where to get your content when users first request it and also when it needs to be refreshed in the cache	Enter the domain name or IP address of your origin server in the **Source hostname** field. Please note that if you enter a domain name, it must be fully qualified.
Source URL path	The specific path of the source hostname that contains content.	The Content Delivery service needs to know the specific path because a source hostname can contain many paths.	If your content is all in a particular path on your origin, or you added a directory name "alias" with the Published Hostname for a particular origin path, you can enter the origin path by clicking the **Add Path** link **Note:** This field must contain a value and defaults to /
Source port	The HTTP port number to use when communicating with your origin server, using the **Origin Host** and **Origin Path** you specified	If you are using a port other than the default (80) for HTTP, the *Content Delivery* service needs to know which port you've chosen	Leave the default port number for HTTP unless you are using another port number. If so, enter the new port number in the **Origin HTTP Port Number** field. **Note:** The default for HTTPS is 443, and this is the value used by Limelight for all HTTP requests to origin (the value is not editable).
Location of Content Source	The location of the content you want the *Content Delivery* service to deliver (the "origin")	The *Content Delivery* service needs to know where to find your content when users first request it and also when it needs to be refreshed in the cache	If your content is not stored with Limelight, choose **Outside Limelight infrastructure**. Otherwise, choose the appropriate Limelight Storage location. If you choose **Outside Limelight infrastructure** in **Location of Content Origin**, the following additional fields are displayed: • Origin Protocol • Origin Host • Origin Path • Origin HTTP Port If you choose a Limelight storage option in

			Location of Content Origin, the following fields are displayed: • Origin Path **Note:** If you are using Limelight storage but your storage option is not shown, your *Content Delivery* service is not fully configured. If this is the case, please contact Limelight Customer Service.

Table 1. Configure - Delivery - Content Location Settings

Host Header Details

Browsers usually include the origin domain name of the requested URL in the HTTP `Host` header. You can use this behavior to detect and block such requests on your origin, denying those with a `Host` header that matches your domain name and accepting those that match either your **Published hostname** or another value you enter in the **Value** field.

If you are hosting more than one origin on a single server and want to block based on Host headers, don't use **Published hostname**; instead, enter a value in the **Value** field. If you are hosting more than one origin on a single server and don't want to block them based on `Host` headers, choose **Origin Host**.

Example Settings

Configuration Field	Value	Notes
Published protocol	HTTPS	Accept only HTTPS requests for cached content
Published hostname	`published.host.com`	Use a CNAME alias instead of the name provided in the Welcome Letter (need to set up the CNAME separately)
Published URL path	`/pubimages/`	Use the pubimages directory to identify the content in cache uniquely.
Source protocol	HTTP Always	Always use HTTP to communicate with the origin server
Source hostname	`origin.host.com`	
Source URL path	`/images/`	Directory path to the origin content; note that this doesn't need to match the path (if any) for the Published Hostname
Origin HTTP Port	`80`	Use the default HTTP port (no need to change anything)
Host Header	Published Hostname	Blocks most browser requests made directly to the origin

Table 2. Configure - Delivery - Content Location - Example Settings

Using the example configuration settings above, if `favicon.ico` is not cached for this configuration or has expired in the cache, a request to `https://published.host.com/pubimages/favicon.ico` will result in an origin request for `http://origin.host.com/images/favicon.ico`, with an HTTP `Host` header of `published.host.com`.

Caching Rules

Setting	Information Requested	Purpose	Selecting the Right Option
TTL	Whether to override the default method for determining if an object in the cache is expired.	In some cases, you may want to take explicit control over object expiration times (TTL - "Time To Live").	Make selections in the **TTL presets** field. • To allow *Content Delivery* to calculate TTL, select **Not selected**. • To make your TTL settings, select **Configure manually** and configure the **Minimum TTL** and **Maximum TTL** fields. When you make manual settings, you must also configure the **Cache Generated Responses** field. • To set TTL to a specific length of time, select one of the preset times. When you opt for preset values, you must also configure the **Include generated responses** and **Cache Generated Responses** fields. See Caching Best Practices for additional information.
Cache large files on first request	Whether you want any request for an object to force the full object to be cached, even if the request is canceled.		**Note:** This feature is intended for large file downloads and is not recommended for caching website objects (such as image, CSS, and JavaScript files).
Ignore "No cache" header	Whether *Content Delivery* should ignore certain `Cache-Control` headers when determining whether or not to cache an object retrieved from your origin	You may want to cache objects regardless of origin settings that attempt to turn off caching	If you want to ignore the following `Cache-Control` headers: • `Cache-control: no-cache` • `Cache-control: no-store` • `Cache-control: private` • `Pragma: no-cache` enable this option. Otherwise, leave it disabled.
Query String Caching	Whether to use URL query terms to determine whether or not objects are cached	You may want to increase cache efficiency by ensuring certain objects are not duplicated due to variations in their query terms	Choose the option that caches the minimum number of objects necessary based on query parameters: • **Strip no query terms from the cache key** • **Strip all query terms from the cache key** • **Exclude specific query terms** • **Keep only specific query terms** **Note:** For the **Exclude specific query terms** and **Keep only specific query terms** options, you must enter a comma-separated

			list of the query terms to be excluded or included
Vary Headers **Note:** for static content configurations only	Which `Vary` response header fields Content Delivery should use when differentiating versions of an object in the cache	Content Delivery stores a different version of a requested object for each unique set of request header fields specified by the `Vary` header. If the `Vary` header specifies request header fields that change frequently, multiple copies of the same object may be stored in the cache. To control this behavior, you can configure *Content Delivery* to ignore all `Vary` headers or specific `Vary` headers when caching and retrieving objects. All of the Vary headers associated with the object are still maintained and passed on to the client in the response.	• If you only want to cache a single version of an object regardless of its `Vary` header fields, choose **Ignore all Vary headers** • If you want to cache a new version of an object whenever any of its `Vary` header fields changes, choose **Do not ignore Vary headers** • If want cache a new version of an object whenever all but certain specified `Vary` header fields change, choose **Ignore specific vary headers** and select the `Vary` headers fields to ignore
Partial Cache	Whether to use Partial Caching to improve cache performance	Partial Caching is a *Content Delivery* feature that caches commonly-requested portions of content requested using HTTP GET ranges. This optimization can significantly improve performance for large media files.	To enable this setting, check the **Partial Cache** checkbox, and in the associated field, enter a Regex value that matches the object URLs you want to optimize
N Byte Download	Whether to download the first "n" bytes to improve cache performance	*Content Delivery* can automatically cache a specified number of bytes from the beginning of cached files. This optimization can improve first-byte response times in some scenarios.	To enable this setting, check the **N Byte Download** checkbox, and in the associated field, enter a Regex value that matches the object URLs you want to optimize

Table 3. Configure - Delivery - Caching Rules Settings

Caching Best Practices

Use the **Not selected** setting unless you have a specific reason to override how the Time to Live (TTL) is calculated.

By default, *Content Delivery* considers an object “stale” (expired from cache) if the number of seconds specified by the associated `Cache-Control: s-maxage` or `Cache-Control: max-age` header has elapsed since initial caching or since the last freshness check, or if neither header is present if the date and time in the `Expires` header have passed. The order of precedence is `Cache-Control: s-maxage`, `Cache-Control: maxage`, then `Expires`.

If no explicit freshness information is supplied (there are no `Cache-Control: s-maxage`, `Cache-Control: max-age` or `Expires` headers), and a `Last-Modified` header is present, the CDN will by default use the adaptive cache freshness algorithm to calculate remaining TTL, based on 20% of the age of the cached response, subject to a floor of 3 seconds and a ceiling of 3 days.

If you need to override (ignore) the above behavior, you can use the **Configure manually** option to specify a new TTL value using the **Time to live (TTL)**drop-down menu.

You can also control whether generated responses are cached using **the Cache Generated Responses** checkbox (for the **Custom** option) or **Including Generated Responses** (for other values in the drop-down menu).

Notes:
Generated responses are HTTP responses that are generated dynamically ("dynamic content"). These responses often do not include any of the cache-control headers needed to determine TTL and are not cached by default to avoid caching personalized or user-specific responses.

By default, Limelight defines a generated response as one that is missing all of the following headers:

- `Expires`
- `Last-Modified`
- `Cache-Control: max-age`
- `Cache-Control: s-max-age`

If you choose the **Custom** option for **Time to live (TTL)**, you can change the cache freshness algorithm's parameters using **Specify custom floor and ceiling cache values**.

If desired, the floor (minimum) can be raised, and the ceiling (maximum) can be lowered or raised. If min and max are set equal to each other, the TTL becomes explicit rather than adaptive.

Arc Light

You can use Arc Light to customize how Content Delivery reacts to HTTP requests and responses. Rules can be triggered when a request or response meets pre-defined conditions, such as a pattern match. Rules are designed based on specific customer needs.

Note: This option is available only for *websites & apps* configurations

Configuration Overview

Use Arc Light to customize how Content Delivery reacts to the following HTTP request and response types:

- Requests
  - Any
  - Origin only
  - Edge only
- Responses
  - Any
  - Origin only
  - Client only

For each of the above request and response types, you can assign one rule. Content Delivery will then execute that rule each time it receives the associated request or response type. The rule will be executed on the Edge Server that receives the request or response.

To enable Arc Light for a specific request or response type, check the checkbox next to the desired type (example: ☒ **Rules on Edge Request**). To assign a rule, click one of the rules in the list below the request/response type.

Rules can be triggered when a request or response meets pre-defined conditions, such as a pattern match with:

- The URL, file name, or query term
- The IP address
- The value of a specified HTTP header
- A cookie
- The geographic location of a request (using the IP address)

When a rule is triggered, it can perform a variety of actions, such as:

- Controlling which CORS headers are sent in response to a client request
- Adding a cookie that contains geolocation information
- Adding specific HTTP headers
- Appending special “keys” to cache keys Enabling or disabling GZIP compression
- Controlling whether the requested content is cached and setting content TTLs

Rules are designed based on specific customer needs. If you need to use Arc Light or want more information on the types of rules that can be created, please contact your Account Manager or Solutions Engineer.

Configuration Settings

Setting	Information Requested	Purpose	Selecting the Right Option
Which rules do you want to enable?	If you want to create a new rule, the type of HTTP request or response to associate it with	Content Delivery can trigger rules for several types of requests and responses	(see the options below)
Rules on Any Request	Request type	Content Delivery can trigger rules for several types of requests	To trigger a rule on any request received by a LimelightEdge Server, check the **Rules on Any Request** checkbox, and select one of the predefined rules in the list
Rules on Edge Request	Request type	Content Delivery can trigger rules for several types of requests	To trigger a rule on client requests to a LimelightEdge Server, check the **Rules on Edge Request** checkbox, and select one of the predefined rules in the list
Rules on Origin Request	Request type	Content Delivery can trigger rules for several types of requests	To trigger a rule on Limelight requests to your Origin, check the **Rules on Origin Request** checkbox, and select one of the predefined rules in the list
Rules on Any Response	Response type	Content Delivery can trigger rules for several types of responses	To trigger a rule on any response received by a LimelightEdge Server, check the **Rules on Any Response** checkbox, and select one of the predefined rules in the list
Rules on Origin Response	Response type	Content Delivery can trigger rules for several types of responses	To trigger a rule on responses received from your Origin, check the **Rules on Origin Response** checkbox, and select one of the predefined rules in the list
Rules on Client Response	Response type	Content Delivery can trigger rules for several types of	To trigger a rule on responses received from the requesting client, check the **Rules on Client Response** checkbox, and select one

	responses	of the predefined rules in the list	

Media Delivery

Content Delivery supports "seeking" or "scrubbing" (skipping back and forth) within FLV and MP4/H.264 video files. Seeking is controlled via parameters specified in the query terms of the request URL.

Setting	Information Requested	Purpose	Selecting the Right Option
Enable FLV Scrubbing	Whether to allow a video client to skip forward and back (seek) within FLV files based on parameters specified in the query terms of the request URL.	Custom clients may want to provide the “seek” capability (“forward” and “back” buttons)	To enable this feature, check the **Enable FLV Scrubbing** checkbox
Enable MP4/H.264 Scrubbing	Whether to allow a video client to skip forward and back (seek) within MP4 files based on parameters specified in the query terms of the request URL.	Custom clients may want to provide the “seek” capability (“forward” and “back” buttons)	To enable this feature, check the **Enable MP4/H.264 Scrubbing** checkbox

Table 4. Configure - Delivery - Media Delivery Settings

Optimization

Setting	Information Requested	Purpose	Selecting the Right Option
Type of Compression	Whether to use Gzip compression when delivering XHTML, JavaScript, CSS, and other text files	Compressed objects are delivered more quickly, potentially improving the user experience	• If you want to provide all compressed files from your origin server, choose the **Gzip Passthrough** option • If you prefer to have the *Content Delivery* service compress files when the requesting client can accept them, choose **Gzip on-the-fly** • If you need to modify Gzip compression defaults, choose **Custom**, then either **Gzip on-the-fly** or **Gzip Passthrough**, and enter your Gzip modification extensions • You can also choose **No compression** if none of your files should be delivered compressed • For more information on this feature, see Gzip Details
TCP Acceleration	The “profile” to use when accelerating the transfer of IP packets by modifying default TCP parameters	In certain circumstances, you may want to change the **TCP Acceleration** profile to optimize your delivery performance	When **TCP Acceleration** is enabled, the XDLL profile is the most efficient in many cases. **Note: TCP Acceleration** is an advanced configuration setting and should only be changed if you’re an expert user.
Enable chunked response to client	Whether the *Content Delivery* service can maintain open TCP connections to your	If there is a cache “miss” and your origin doesn’t provide a `Content-Length` header, this option allows *Content Delivery* to serve the requested content more efficiently	We recommend you enable this option. Otherwise, new TCP sessions must be established for each new request to origin, and cache miss requests are delivered only when the entire object has been transferred from

Note: for static content configurations only	origin server	(in “chunks”)	origin.

Table 5. Configure - Delivery - Optimization Settings

Gzip Details

When**Gzip Passthrough** is enabled, and a client indicates (via HTTP request header) that it prefers to receive compressed content, the*Content Delivery* service will serve a compressed version of the requested object if one is available on the origin server.

> **Note:** ***Gzip Passthrough*** *is available to all customers. If it is not enabled for you, please contact Limelight Support.*

If**Gzip On-the-fly** is selected, the*Content Delivery* service creates, caches, and delivers Gzip-compressed content as needed.

Compressible file types include:`action, ashx, asmx, asp, aspx, axd, cfm, css, css3, csv, do, doc, docx, htm, html, js, jsf, json, jsp, php, portal, rtf, svg, svgz, tsv, txt, xhtml, xml, site root (/),` and extensionless URLs.

Headers & Methods

Setting	Information Requested	Purpose	Selecting the Right Option
Client Analytics	Whether you want the *Content Delivery* service to provide geographic user information when requesting content from your origin	You may want to capture, analyze, and report on user geographic information internally.	To use this feature, check the **Client Analytics** checkbox. The geo information is provided to your origin server via two request headers: `X-IP-Geo-Country` and `X-IP-Geo-All`. The geo fields provided are `continent`, `state`, `city`, `dma_id`, and `asn`.
Add client IP address to origin request header **Note:** for static content configurations only	Whether you want the*Content Delivery* service to provide the requesting client's IP address in a custom header when requesting content from your origin	You may want to capture, analyze, and report on user IP information internally.	To enable this feature, check the**Add client IP address to origin request header** checkbox, and enter the header names containing the client IP address. The default header name is`True-Client-IP`. Note that the above headers are in addition to`X-Forwarded-For`, which is always provided to the origin.
POST Requests	Whether you want to accept or ignore POST requests from clients	If you are using a custom client to display content, you may want to communicate analytics or other information to your origin. Alternatively, you may want to convert POST requests to GET requests or ignore them.	• To ignore all POST requests, select **Disable HTTP POST requests**.*Content Delivery* will respond with an HTTP `413 Request Entity Too Large` status code to all POST requests.

			• To accept POST requests and pass them through to your origin, select **Enable HTTP POST requests**. If a POST request body exceeds 500 MB, *Content Delivery* will respond with a `413 Request Entity Too Large` status code. • To accept POST requests but treat them as GET requests, select **Enable HTTP POST requests,** and check the **Discard request body on POST request** checkbox. POST bodies will be discarded.
Add custom request header	Whether you want to include custom headers and values whenever *Content Delivery* makes a request to your origin	If you want to tag all requests from *Content Delivery* for later analysis	To add a custom origin request header, enter a unique header name and value and. Click the "+" button to add additional headers.
Add Limelight server IP address when responding to client	Whether to provide clients with the IP address of the *Content Delivery* Edge Server responding to their requests	If you are using a custom client to display content, and you are also capturing performance-related data via the client, you may want to include the *Content Delivery* Edge Server IP address for later analysis and reporting. The IP address will be provided in the X-IP-Address response header.	To enable this feature, check the **Add Limelight server IP address when responding to client** checkbox
Add custom response header	Whether you want to include custom headers and values whenever *Content Delivery* responds to a client request	If you are using a custom client to display content, you may want to provide it with information that uniquely identifies the *Content Delivery* service, Limelight Account, etc.	To add a custom client response header, enter a unique header name and value and. Click the "+" button to add additional headers.
Enable Custom Debug Headers	Whether you want to enable Custom Debug Headers	By making an HTTP content request with special "Custom Debug Headers," including a shared secret specific to your service, you can retrieve cache-related information about individual content objects and prevent others from accessing the information.	In the **Debug Headers** field, enter one or more "tags" to include in the Custom Debug Headers. Then in the **Secret Key To Request Debug Information** field, enter the secret key (shared secret) provided by Limelight when the Custom Debug Headers feature was enabled. For more information, see Secure Cache Diagnostics.

Table 6. Configure - Delivery - Headers & Methods Settings

Secure Cache Diagnostics

When troubleshooting caching issues, customers can directly access diagnostic information about cached content.

Content Delivery customers with *Configuration Self Service* can request cache-related information about individual content objects and prevent others from accessing the information.

To enable this feature, check the *Enable Custom Debug Headers* checkbox in the *Request and Response Headers* section of *Content DeliveryConfiguration Self Service*, and provide a comma-separated list of object properties that should be returned in the response, along with a Secret Key to authenticate the request.

Diagnostic response headers can include the following information:

- Whether or not a response is cacheable
- How the cache responded to a request (hit, miss, etc.)
- The number of seconds before the cached response will be considered stale (TTL)
- The total number of seconds representing the freshness lifetime of the response (age + TTL) and how the value was determined (headers, overrides, adaptive TTL, etc .)

When the feature is activated, you will be provided with a unique shared secret.

The properties that can be requested, and their associated response headers and values, are:

Request Key	Response Header	Return Values

`is-cacheable`	`X-LLNW-Dbg-Is-Cacheable`	`Yes`
		`No`
		`Negative`

`cache-hit-type`	`X-LLNW-Dbg-Cache-Hit-Type`	`HIT`
		`MISS`
		`REFRESH_HIT`
		`REF_FAIL_HIT`
		`REFRESH_MISS`
		`CLIENT_REFRESH_MISS`
		`IMS_HIT`
		`NEGATIVE_HIT`
		`DENIED`
		`OFFLINE_HIT`
		`REDIRECT`

`ttl`	`X-LLNW-Dbg-TTL`	`n{...} seconds` an integer followed by a space and the string "seconds"
`fresh-life-total`	`X-LLNW-Dbg-Fresh-Life-Total`	`n{...} seconds` an integer followed by a space and the string "seconds"

If the secret is invalid, the `X-LLNW-Dbg-Hdrs` header will be ignored, and the request will be processed without it.

Request and response example:

Request	Response
`GET http://www.customer.com/object.txt HTTP/1.1...X-LLNW-Dbg-Hdrs: is-cacheable,cache-hit-typeX-LLNW-Dbg-Secret: sharedsecret`	`HTTP/1.1 200 OK...X-LLNW-Dbg-Is-Cacheable: Yes...X-LLNW-Dbg-Cache-Hit-Type: HIT`

Failover

Normally when the CDN receives a `404 (Not Found)`, `503 (Service Unavailable)`, or `504 (Gateway Timeout)` response from your origin, the error is passed back to the requesting client. You can modify this configuration option as follows:

- For `404` errors:
  - Serve "stale" content from the CDNCache, or
  - Request content from a "backup host" (with or without a path), or
  - Redirect to a custom "Not Found" URL.
- For `503` and `504` errors:
  - Request content from a "backup host" or
  - Redirect to a custom "Service Unavailable URL."

> **Notes:**
> - Failover URLs must match their configuration within the CDN
> - For `404` error redirects, the original request is reissued to the fallback URL with any modifications still in place
> - `503` or `504` errors may have been generated by the origin but could also be generated by CDN if a connection can't be made to your origin

Setting	Information Requested	Purpose	Selecting the Right Option
Serve stale content instead of 404 error	If the requested content is cached but stale (expired), and there is an HTTP `404` status when requesting a fresh version from your origin, this field indicates whether you want to pass the `404` status back to the client or serve the stale content instead	If an object has expired in the cache and your origin server returns a `404 (Page Not Found)` error when *Content Delivery* attempts to get a fresh copy of the object, you may want to serve the expired object instead of allowing the client to handle `404` messages.	If you don't want the client to handle `404` messages, and it is acceptable to serve stale content instead, check the **Serve stale content instead of 404 error** checkbox. **Note:** If there is no cached object, a 404 message will still be returned to the browser.
Request content from backup host on 404 error	If there is an HTTP `404` status when requesting fresh content from your origin, this field indicates whether to use a backup origin (hostname only) before handling the `404` status	If your primary origin returns a `404` status, and you have a backup origin, you may want *Content Delivery* to try the backup before handling the error.	To serve content from a backup origin if the primary origin responds with a `404` status, enter the fully qualified hostname of the backup origin. **Note:** Specific ports are not supported.
Use custom "Not Found" page	Whether you want to pass HTTP `404` status messages back to the client or serve a custom error page instead	If an object has expired in the cache, and your origin server returns a `404` error to *Content Delivery*, you may want to serve a custom error page rather than allowing the client to handle the `404` message.	To take control over clients' content when the origin returns a 404, enter the fully qualified URL of the content to serve.
Request	If the origin responds with an HTTP 404	If your primary origin returns a `404` status,	To serve content from a

content from backup origin URL on 404 error	status upon request for fresh content, this field indicates whether the request should be sent to a backup URL path before handling the `404` status	and you have a backup origin, you may want *Content Delivery* to try the backup before handling the error.	backup URL path if the primary origin responds with a `404` status, enter the fully qualified path on the backup origin. **Notes:** • You can specify either the HTTP or HTTPS protocol and a port number if desired. • This option is required when using the Intelligent Ingest feature of Origin Storage
Request content from backup host on 5xx error	If there is an HTTP `5xx` status when requesting fresh content from your origin, whether to try a backup origin before handling the `5xx` status	If your primary origin returns a `5xx` status, and you have a backup origin, you may want *Content Delivery* to try the backup before handling the error.	To serve content from a backup origin if the primary origin responds with a `5xx` status, enter the backup origin's fully qualified hostname. **Note:** Specific ports are not supported.
Use custom "Service Unavailable" page	Whether you want to pass HTTP `503` and `504` status messages back to the client or serve a custom error page instead	If an object has expired in the cache, and your origin server returns a `503 Service Unavailable` or `504 Gateway Timeout` error to *Content Delivery*, you may want to serve a custom error page instead of allowing the client to handle the error message.	If you want to take control over the content displayed by clients when there is a `503` or `504` error from the origin, enter the fully qualified URL of the content to serve.

Table 7. Configure - Delivery - Failover Settings

Content Security

IP Access Control

This section allows you to allow or deny access to content based on IP addresses and geographic locations ("geo-fencing").

Content Delivery Configuration Self Service provides access control using IP addresses and geographic locations ("geo-fencing"). When configuring an Account, you can associate lists of IP addresses and groups of geographic locations with the Account and specify whether to allow or deny each. When managing IP address lists, you can also view whether they are currently in use and which Accounts they are associated with (or limited to).

Setting	Information Requested	Purpose	Selecting the Right Option
Enable IP Access Control	Whether you want to "allow list" or "deny list" requests based on IP address lists and IP-based geographic locations	IP Access Control allows you to exclude specific geographies or limit access to known entities	Assign access lists to the Caching & Delivery configuration using the following drop-down menus: • **By IP address list**: Select one or more existing lists, then choose either **Deny** or **Allow** to indicate the type of restriction. Click **Add** to add the lists to the **Access control list for this configuration** section. • **By geolocation**: Select one or more geographic areas (continents or countries), then choose either **Deny** or **Allow** to indicate the type of restriction. Click **Add** to add the lists to the "Access control list for this configuration" section. **Access control list for this configuration** Section You can select a default security setting for the configuration - either **Default Allow** or **Default Deny**. You can then add one or more IP address lists and geographic locations that modify the default setting. IP address lists and geolocations can be "mixed and matched" in any order desired. To move an item in the list, move the mouse pointer over the item and use the vertical ellipses to drag and drop the item to another location in the list. If you have the correct permissions, click **Manage IP Lists** to display a dialog that allows you to create new IP access lists. You can also view, edit, and delete existing IP address lists. To view list details, click the **+** icon to the left of a list. • The text "Used by configs in accounts" shows which Accounts have configurations that use the list. • The text "Limited to accounts" shows any accounts to which your *Company Admin* has limited the list. To create a new list, click the **new list** button at the top of the dialog, then: 1. Provide a name for the list. 2. Provide a single IP address or range of IP addresses. You can also create and upload CSV files of IP addresses. Click the link to see a sample CSV file. 3. Optionally limit the list to accounts. 4. Click the **Save** button. The new list is now available in the **By IP address list**: drop-down menu at the top of the section. To deny access to end users attempting to access your content using an anonymous VPN from an unauthorized geolocation. select the **Deny 'Anonymized with VPN' access** option. **Notes:** • IP address lists and geographic locations are processed in the order they are specified (top to bottom). Once a match is found, subsequent lists and locations are ignored • Users with the *Company Admin* role can manage lists for all accounts. Users with the *User* role who have been granted the *Manage Delivery Configurations* permission can apply all lists in the Accounts for which they have been granted management permission. • Changes made to IP address lists are applied immediately and affect all Account configurations which use them (even legacy configurations that can't be edited in Control. • IP address lists cannot be deleted if they are in use.

Table 8. Configure - Delivery - IP Access Control

MediaVault

Setting	Information Requested	Purpose	Selecting the Right Option
Enable MediaVault content protection	Whether you want to use *MediaVault* to provide additional content security. *MediaVault* provides high-performance URL authentication.	*MediaVault* can help you prevent "deep linking" and other unauthorized viewing behavior	To enable this feature, check the **Enable *MediaVault* content protection** checkbox, and provide a primary and secondary "shared secret" (both used to prevent URL tampering). You can also change the HTTP Error Code returned by MediaVault from the default 400 code by entering a new value in the **Deny Status Code** field. For more information, see *MediaVault* Details.

Table 9. Configure - Delivery - IP Access Control

More about MediaVault

MediaVault is a high-performance URL authentication service. *MediaVault* 's main purpose is to help you secure your content from unauthorized viewing.

MediaVault maximizes authentication performance by using tokens to avoid three-way handshakes (common to other authentication methods) that can lead to severe connection time latency.

Please note that *MediaVault* is *not* a replacement for DRM and should not be associated with user authentication.

MediaVault works like this:

- You enter a shared secret during the configuration process
- You then generate a token (MD5 hash) for each published URL, based on the shared secret, and append it to the URL in a query term or provide it in a cookie.You can generate the token manually by navigating to the *Configure* > *MediaVault* in the navigation pane, or by creating server-side code on your origin.
- *MediaVault* uses the same hash algorithm to create its token when a request is received, identical to the one you appended.
- If the tokens match, *MediaVault* then looks for additional *MediaVault* -specific query terms (such as end date/time and IP address/mask) to determine whether the request is valid. If the tokens don't match, the URL was tampered with, and the request is rejected.

For more information, see the MediaVault User Guide by navigating to Help Center > Documentation > Delivery > Guides > MediaVault in the navigation pane.

> **Note:** *MediaVault* does not support IPv6-based Client IP or IP Range restrictions at this time.

Amazon S3 Authorization

If you store content on Amazon S3, use this option to set your S3 access key, secret, and region.

Send SSL SNI to Origin

Server Name Indication (SNI) is a TLS extension that allows multi-tenancy of domains hosted on a web server. Shared cloud platforms often require SNI. The extension helps select the appropriate certificate for that domain and helps serve the appropriate content. Most modern web servers handle SNI; this option disables SNI for the minority of web servers that do not handle SNI correctly.

- If a user is creating a new configuration and the selected protocol sets include HTTPS, then the **Send SSL SNI to Origin** option is selected by default.
- If a user is editing a configuration, then the field is visible and enabled depending on the selected Service Profile.

Logging

Setting	Information Requested	Purpose	Selecting the Right Option
Log cookies	Whether you want *Content Delivery* to *stop* saving cookie information in your log files	If you process log files and don't need the information in the Cookie header, you may want to remove it to simplify processing and reduce log file size.	If you know you need Cookie header information in your log files, check the **Log cookies** checkbox. Otherwise, leave it unchecked. When this setting is enabled, *Content Delivery* logs all `Cookie` header information, up to a maximum of 8 KB for the entire header (regardless of the number of cookies in the header).
Log request header	Whether you want *Content Delivery* to *start* saving specific Request Headers in your log files	If you process log files and need access to information in the Request Headers, you may want to enable this option	If you know you need Request Header information in your log files, check the **Log Request Header** checkbox and enter the names of the specific headers to log. Otherwise, leave it unchecked.

Table 10. Configure - Delivery - Logging Settings

Cookie handling

EdgePrism issues a `Set-Cookie` header whenever it receives a request that has a specified query parameter. This feature provides a way for a cookie to be set with the query string sent in a request URL. You can create a configuration by entering values into the fields provided.

Field	Description
Cookie parameter	Cookie name.
URL query term selector	String that identifies the query term.
Expiration	Date when the configuration expires.
Domain	Domain to which the configuration applies

Note: You can also use the feature to indicate when a `Set-Cookie` header should not be issued.

Cookie Handling Example

This example instructs EdgePrism to issue a `Set-Cookie` header to the requesting client with the key `nlpqtid`, no expiration, and a Domain parameter of .ExampleDomain.com whenever EdgePrism detects `pid=` in the requested URL's query string.

Field	Value
Cookie parameter	nlpqtid
URL query term selector	pid=

Field	Value
Expiration	0
Domain	.ExampleDomain.com

Redirect

You can specify conditions under which a content request should be redirected.

Redirect hostname header regex

This option can be used to issue a redirect based on a specified header and value. You can optionally request strict header regex checking.

Field	Description/Instructions
Header name	Name of the header on which the redirect is based.
HTTP Code	HTTP Status Code upon which to issue the redirect.
Comma-separated key-value pairs	Header values upon which to issue the redirect.
HTTP code 301 or 302	Status code to use for the redirect. Enter either 301 or 302. Instead of delivering content from the origin, EdgePrism can redirect the user to a particular URL.

Strict header regex checking

If the header values you entered in the *Redirect hostname header regex* are not in the specified Header name, EdgePrism will use the fields below as the conditions under which it will issue the redirect.

Field	Description/Instructions
HTTP Code	HTTP status code returned from the content request.
hostname	Hostname from which content was requested.

Others

This section presents additional delivery options you can use in the Chunked Streaming configuration. For descriptions, hover your mouse pointer over the right side of the option name. An information icon appears along with the option description.

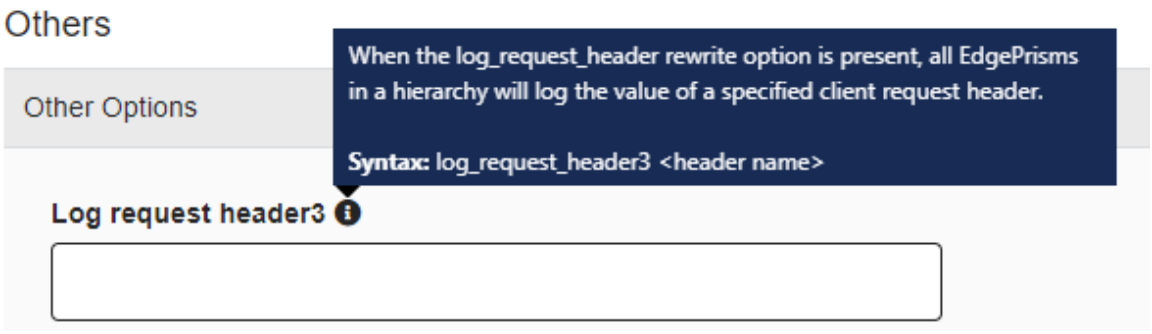

Additional Options

The **Additional Options** section allows you to quickly configure options that are available elsewhere on the page. If you know the options you want, you can configure them here in one location.

If you selected **Both HTTP and HTTPS** in the Content Location section, you could use the section to configure options for a particular protocol set.

1. Begin typing an option name in the Options field. The field has auto-complete capabilities, so you do not have to type the full name. As you type, matching options display in the auto-filtered list.

   > **Note:** Available options depend on your account name and the service profile at the top of the page.

2. Select the option from the auto-filtered list.
   The UI adds the option to a list above the option field. See Working with the Options/Protocol Sets List.

Working with the Option/Protocol Sets List

This list allows you to associate protocol sets for the option you selected and enter any required option parameters.

1. If parameters are required for the option, a field is displayed to enter parameter values. A prompt describing the parameter is displayed beneath the field.
2. Enter a parameter value.
3. If you selected **Both HTTP and HTTPS** in the Content Location section, two protocol sets are displayed to the right of the list item; otherwise, only one protocol set is displayed.
4. Associate an option set with the option by selecting the desired protocol sets.
5. To remove an option, hover over its row in the list and click the **(x)** icon on the right side of the row.

Notes

You can use the **Notes** field for additional information for others (why the configuration changes were made, etc.). Users can refer to the notes later when browsing historical configuration changes

Editing a Configuration

To edit a configuration, click the "pencil" icon at the bottom right of the configuration row. For information on the individual settings displayed, please see the descriptions in Creating a New Configuration.

Notes:

- The ability to edit configurations is subject to conditions described in Read-Only and Hidden Capabilities.
- On rare occasions, a configuration might contain unsupported protocol set configurations, and if you attempt to edit the configuration, Control prevents you from editing and displays this message:
  "This protocol combination is not supported in this application. You may view it here, but to make modifications, either do so via our configuration API or reach out to your account team."
- Unsupported protocol sets are generally the byproduct of migrating a configuration from an older configuration version.

If you want to change protocol sets, see Changing Protocol Sets.

Note: For some configurations created for you by Limelight, only the preview ("eye") icon will be visible. If this is the case, and you need to make changes, please contact your Account Manager or Solutions Engineer. Limelight can continue to manage the configuration, or it can be made available for you to edit in the *Limelight Control*.

Previewing a Configuration

To preview the settings associated with a configuration, click the "eye" icon at the bottom right of the configuration row. For information on the individual settings displayed, please see the setting descriptions in Creating a New Configuration.

Cloning a Configuration

To clone (make a copy of) a configuration, click the "copy" icon at the configuration row's bottom right.

When you have finished making changes to the settings, click **Activate** to enable the new configuration.

If you want to change protocol sets, see Changing Protocol Sets.

Note: The ability to clone configurations is subject to conditions described in Read-Only and Hidden Capabilities.

Note: For some configurations created for you by Limelight, only the preview ("eye") icon will be visible. If this is the case, and you need to make changes, please contact your Account Manager or Solutions Engineer. Limelight can continue to manage the configuration, or it can be made available for you to edit in the *Limelight Control*.

Deleting a Configuration

Note: Users are not able to delete configurations. For more information, see Read-Only and Hidden Capabilities.

Reverting to a Previous Configuration

Each time you update a configuration, a new version is assigned.

To revert to a previous configuration:

1. Click the "undo" icon at the bottom right of the configuration row.
   A list of previous versions is displayed in a dialog.
2. Select the version to which you want to revert.

   Note: Although you intend to revert to a previous version, the reverted version will become the current version with a new version number. The new version number is displayed at the bottom of the dialog.

3. Click the **Activate** button.

Note: The ability to revert configurations is subject to conditions described in Read-Only and Hidden Capabilities.

Configuring Intelligent Ingest

If you're an Origin Storage (previously *Cloud Storage*) customer, you can use Intelligent Ingest to automatically populate Origin Storage with new content as it is requested from the CDN. If a request results in a CDN Cache Miss, and the content is not found in Origin Storage, Intelligent Ingest will retrieve and ingest the content from the remote host(s) you specify. You can create any number of "rules" that pair specific Origin Storage logins and content paths with specific remote hosts and paths. These rules are also known as "rewrites."

Intelligent Ingest List Page

The Intelligent Ingest list page provides summary information at the top, including **Status** (*Active*, *Inactive* or *Disabled*), **Storage Quota** (if quota is set; % and total disk usage), **Bandwidth** (the ingest bandwidth limit for the remote hosts), and **Concurrency** (maximum number of concurrent connections to the remote hosts).

A list of existing Intelligent Ingest rules appears below the summary area. For each rule, the content paths for both Origin Storage and the remote host are shown along with the authentication type, if any.

Notes:
- By default, the Intelligent Ingest configuration page is accessible only to internal Limelight users.
- The data for Storage Quota, Bandwidth, and Concurrency was provided by your company when Intelligent Ingest was purchased.
- If the Intelligent Ingest service has become inactive because the quota has been exceeded, you will see an *Inactive* alert in the **Status** section of the overview. When the problem is resolved, the display will update to *Active*.
- If you're unable to use Intelligent Ingest because the quota has been exceeded, please contact your Company Admin.

Enabling Intelligent Ingest

To enable Intelligent Ingest, please contact your Account Manager. After Intelligent Ingest has been turned up for your account, "Intelligent Ingest" will appear under the Configure menu in the navigation pane.

Note: Intelligent Ingest configuration requires at least one active Origin Storage login.

Configuration Overview

To successfully configure Intelligent Ingest, the following steps must be completed:

1. Obtain a Limelight Origin Storage account with an active login. An Origin Storage account is required for an Intelligent Ingest configuration.
2. Enable Intelligent Ingest for that Origin Storage account (see Enabling Intelligent Ingest).
3. Request a Caching & Delivery configuration (or use an existing one) to enable CDN delivery of objects from the Origin Storage account.
4. Create one or more Intelligent Ingest rules (see Adding Intelligent Ingest Rules).

Creating or Editing a Content Delivery Configuration for Intelligent Ingest

As part of Intelligent Ingest configuration, you must create a related Content Delivery configuration.

You can either edit an existing Content Delivery configuration or create a new one.

To work with Content Delivery configurations, perform the following steps:

1. Select *Caching & Delivery (v2)* in the *Configure* tab.
2. Perform instructions in the following sections:
    - Step 1-Configuring Content Location
    - Step 2-Configuring Failover to Backup Origin

Step 1-Configuring Content Location

In the *Content Location* section, make these selections in the *Source* subsection:

1. In the **Location of Content Source** drop-down menu (1), choose the descriptive name that was configured for you. The actual name is displayed in the Source hostname field.
2. Populate the **Base Path** (2) and **Extended Path** (3) fields as needed.

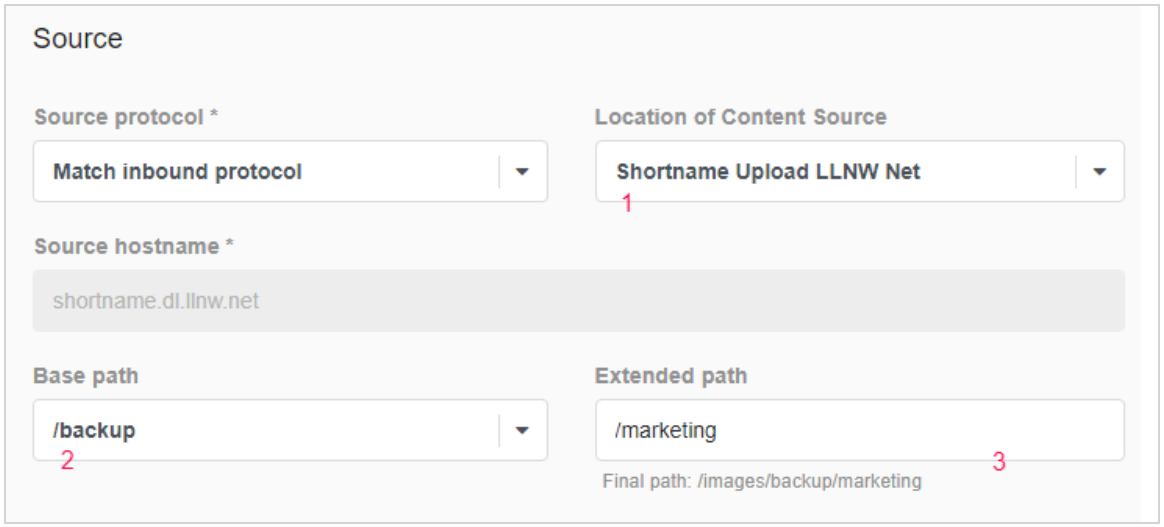

Figure 1. Configure - Intelligent Ingest - Base & Extended Paths

Purpose of Content Location Configurations

Intelligent Ingest rules rely on matching the path of the content you requested from our Origin Storage platform. Generally, the full path on Origin Storage is not something you would have complete information about. The goal is to illustrate the full path of the content location within Origin Storage. The Base path is the part of your Origin Storage path that Limelight creates for you. Rather, it is a drop-down menu, because depending on your permissions, you may have multiple Origin Storage users with various levels of access. Limelight Control then allows you to add an extended path if you don't want to trigger Intelligent Ingest on everything in the base path. The extended path allows more flexibility to use the feature on a subset of content.

Step 2-Configuring Failover to Backup Origin

The Failover to Backup Origin option is required, and must be requested from your Account Manager. Once added, it will be visible in the Content Delivery Configuration Self Service settings:

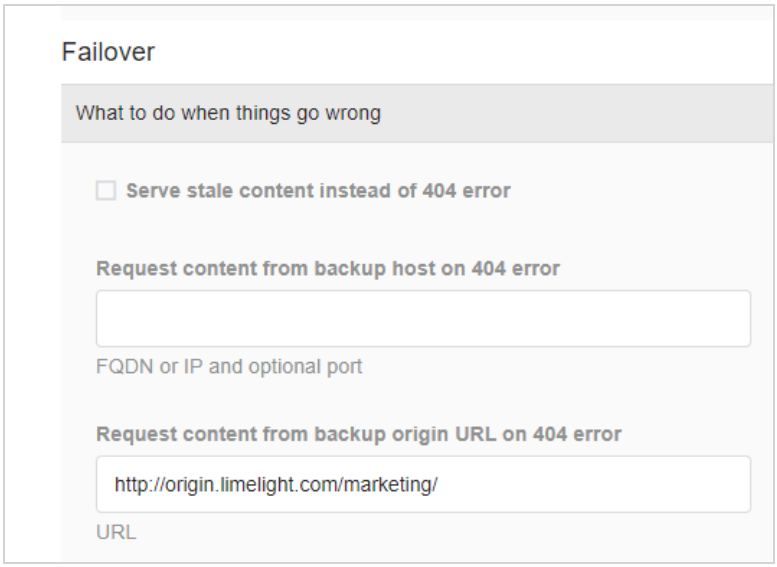

Figure 2. Configure - Intelligent Ingest - Failover Option

Notes:

- You can set up Intelligent Ingest rules and their associated Content Delivery configuration in any order, but both are required before Intelligent Ingest begins ingesting content.
- Although Intelligent Ingest rules require an associated Content Delivery configuration, deleting the configuration does not delete the rules, and deleting rules does not affect the configuration.
- You can configure rules so that content is ingested only for specific paths within a Content Delivery configuration.
- If you need a more complex configuration, more options for your 404 backup origin, or authentication for your backup origin , please contact your Account Manager.

Adding Intelligent Ingest Rules

Before you add a rule, become familiar with how rules are chosen at ingest time. See Workflow Rule Selection.

1. On the Intelligent Ingest page, click the **+ New** button above the summary area. The *Create new rewrite rule* appears with the fields necessary to create the rule.
2. Perform instructions in the following sections:
   - Step 1-Configuring Origin Storage
   - Step 2-Configuring Remote Storage Host
   - Step 3-Testing Paths
   - Step 4-Enabling Remote Storage Host Authentication
   - Step 5-Saving the Intelligent Ingest Rule

Note: At least one rule is required before Intelligent Ingest will begin ingesting remote content.

Step 1-Configuring Origin Storage Origin

The Origin Storage logins associated with the current Account are displayed in the **Origin Storage login** drop-down menu. Select the appropriate user, and also enter the target **Content path** (the path to the Origin Storage directory from which content will be requested, and to which it will be automatically ingested if

requested but not found in that directory).

> **Note:** When you configured the Content Location in *Caching & Delivery (v2)*, you selected the **Base Path** and entered the **Extended Path** (2). The target **Content path** should match the concatenation of those two fields, as shown in Testing Paths below.

Step 2-Configuring Remote Storage Host

For the remote host, under **Remote Protocol**, specify the protocol (*HTTP* or *HTTPS*) Intelligent Ingest should use when requesting your content. Then enter the hostname in **Remote storage host**, and the folder path, if any, in **Remote path**.

Step 3-Testing Paths

To view and confirm the final origin and remote paths, click the **Test** button and review the values in the **Path in Origin Storage** and **Matched remote path** fields.

> **Notes:**
> - The **Test** button is enabled only when the **Path in Origin Storage** field is filled in.
> - Testing is intended to confirm the entered paths. It does not request content or confirm its availability.

The **Path in Origin Storage** field should contain the full path to the Origin Storage directory from which content will be requested, and the **Matched remote path** field should contain the full path to the remote host directory from which the content will be ingested if not found in that Origin Storage directory. If either path is not what you expected, change the values entered for **Content path**, **Remote storage Host** and/or **Remote path** as needed, and repeat the test.

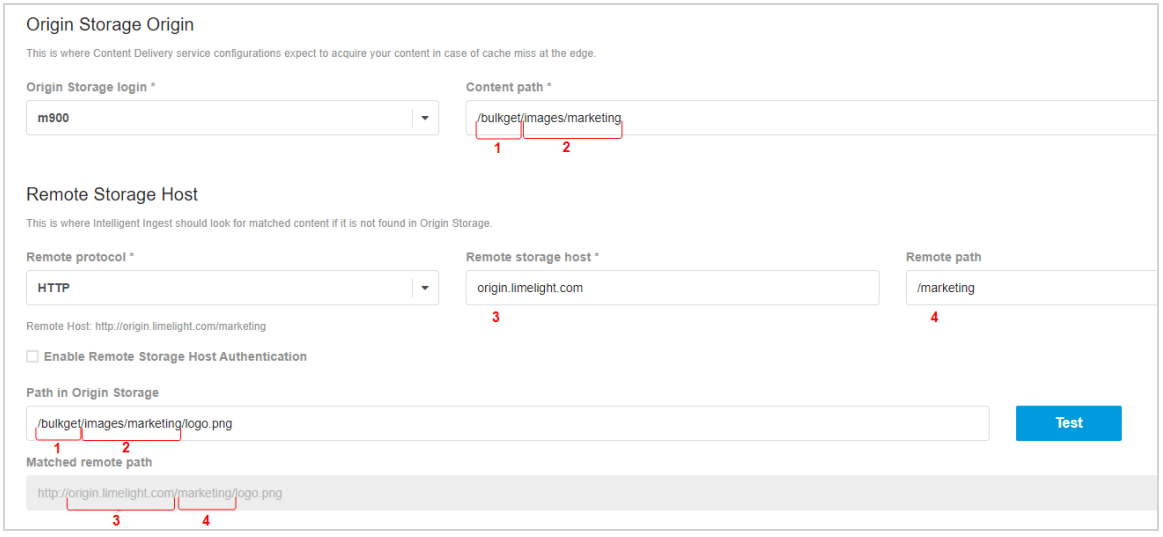

Figure 3. Configure - Intelligent Ingest - Field Comparison

Step 4-Enabling Remote Storage Host Authentication

If your content is hosted on a third-party content provider, you must configure the credentials needed for Intelligent Ingest to access your data. You can choose from existing configurations (1) or create a new configuration (2).

Figure 4. Configure - Intelligent Ingest - Remote Storage Host Authentication

1. Click the **Enable Remote Storage Host Authentication** checkbox.
2. Choose an existing configuration from the drop-down menu or create a new configuration.
3. To create a new configuration:
   a. Click the **+ new** button.
   b. Choose a configuration type in the subsequent dialog.

Type	Description / Instructions
Amazon S3 V4	Select this if your content is hosted on Amazon S3. See Amazon friendly name Fields.
Custom Header	Select this if your content is hosted elsewhere. See Request headers friendly name Fields.

   c. Click **Apply** to create the new rule or **Cancel** to discard your work.

Amazon S3 V4 Fields

This authentication type uses Amazon S3 Signature Version 4 Authentication to send authentication fields to the remote host.

Name	Description / Instructions
Name	Descriptive name of your choice
Access key ID	Amazon AWS Access Key ID
Secret access key	Amazon AWS Secret Access Key
Region	Limelight region into which your S3 content will be ingested. Choose the region that is closest to the S3 region.

Custom Header

This authentication type sends authentication information in request headers to the remote host.

Name	Description / Instructions
Name	Descriptive name of your choice
Add custom request header	1. Enter a name and a value, then click the **+** icon. 2. Repeat until you have entered all required values.

Name	Description / Instructions
	To remove a header, click the **-** icon. After you save the Intelligent Ingest rule (Saving the Intelligent Ingest Rule), the request headers configurations are available to choose for use in future Intelligent Ingest rules.

Step 5-Saving the Intelligent Ingest Rule

When you are ready, click **Save** to create the new rule or **Cancel** to discard your work.

> **Note:** After you save a rule, you cannot modify it; you can only modify any associated authentications (see Managing Authentications).

Managing Authentications

The only component of an existing Intelligent Ingest rule that you can change is its authentications. You can modify, delete, and create authentications.

> **Note:** You cannot restore a deleted authentication.

Begin by clicking the **manage authentication** button on the right above the summary area. Doing so displays a dialog for managing authentications.

To create a new authentication:

1. Click the **+ new** button in the dialog.
   A dialog for creating authentications is displayed.
2. Create the authentication, using instructions in Enabling Remote Storage Host Authentication for creating a new authentication.
3. Click **Apply** to save changes, or **Cancel** to discard your work.

To modify an authentication:

1. Click the **edit** icon for the desired authentication in the dialog.
   A dialog for editing the authentication is displayed.
2. Modify the authentication using information in Amazon S3 V4 Fields and Custom Header Fields.
3. Click **Apply** to save changes, or **Cancel** to discard your work.

To delete an authentication:

1. Click the **delete** icon for the desired authentication in the dialog.
2. Click **Delete** in the confirmation dialog.

Deleting Rules

To delete a rule, click its associated trashcan icon.

> **Note:** Deleting an Intelligent Ingest rule does *not* affect any Content Delivery configuration(s) associated with the deleted rule.

Workflow Rule Selection

Intelligent Ingest includes a workflow in which it attempts to ingest content from origin paths. Understanding the selection process is important before you begin creating rules.

If multiple rules share the same path, the workflow uses path prefix length and order in which rules were created to choose the rule to use.

In the following examples, '/home/data/' is the path prefix.

Example 1

If multiple rules with the same prefix are created in the following order (the shortest is created first), then only the first path will work and the others will fail as part of workflow because they are longer.

1. /home/data/ - works
2. /home/data/dir1 - fails
3. /home/data/dir2 - fails

Example 2

If multiple rules with the same prefix are created in the following order (the shortest is created later), then the first and second paths will work and the third will fail because it is longer than the second.

1. /home/data/dir1 - works
2. /home/data/ - works
3. /home/data/dir2 - fails

Configuring Chunked Streaming (read-only)

Chunked Streaming delivers chunked video content via HTTP and HTTPS in four different formats: HDS, HLS, MSS and MPEG-DASH. To use *Chunked Streaming*, you first need to chunk your content and generate the associated manifest files (*Chunked Streaming* does not perform these operations). You can host your content on your own origin servers or with *Origin Storage*.

> **Note:** Chunked Streaming settings are read-only and will be removed in a future release. Please use Chunked Streaming (v2) instead.

When you select the **Chunked Streaming** option in the **Configure** menu, you'll see the *Configurations* page, with a list of your configurations.

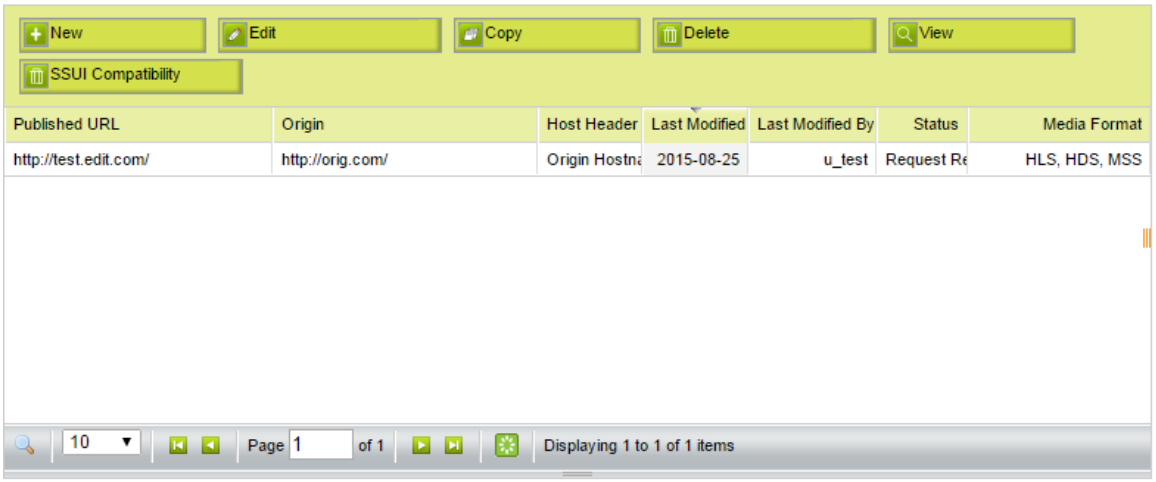

Figure 5. Configurations Page

The *Chunked Streaming Configurations* list provides the following information for each configuration:

- **Published URL** (Published Hostname) – The public URL prefix used in links to your published content (URLs seen by end users)
- **Origin** (Origin Hostname) – The private URL prefix used by Limelight to retrieve and cache content from your origin server (not visible to end users)
- **Host Header Value** – the value that Limelight will include in the HTTP `Host` header when making requests to your origin
- **Last Modified** – The date the configuration was last changed
- **Last Modified By** – The portal ID of the user who last changed the configuration
- **Status** - The processing status of the configuration:
  - **Pending** – The configuration has been submitted, but provisioning has not yet started
  - **In Progress** – The provisioning process has started
  - **Propagating** – The new provisioning configuration is available to all servers that need it, but is not fully deployed
  - **Complete** – The configuration was successfully deployed
  - **Failed** – The configuration was not successfully deployed - please contact Limelight Support for assistance
  - **Unknown** – The status of the configuration could not be determined - please contact Limelight Support for assistance
- **Media Format** - The media formats which the content is delivered in

Creating a New Configuration

To create a new configuration:

- In the *Chunked Streaming Configurations* list, select **New** from the buttons at the top of the list, and the first step of the configuration wizard will be displayed
- Complete each step as necessary, then click **Submit**

Wizard Step: Content Location

In order to fill the cache, the service needs to know where to get your content. In this step, you specify whether you are using CDN Storage or hosting your content outside of the Limelight Infrastructure.

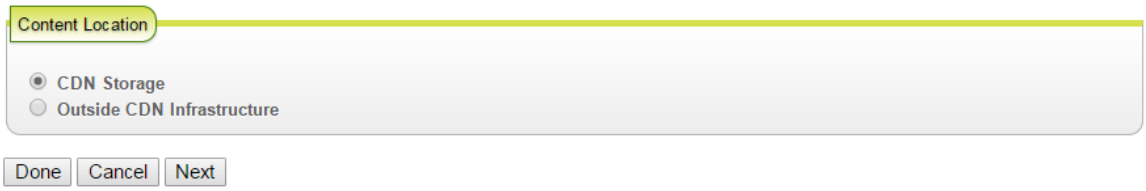

Figure 6. Content Location Step

Setting	Information Requested	Why It's Needed	Selecting the Right Option
Content Location	The location of the content you want the *Chunked Streaming* service to deliver (the "origin")	The *Chunked Streaming* service needs to know where to find your content when users first request it, and also when it needs to be refreshed in the cache	If your content is in *Origin Storage* or *Limelight Discrete Storage*, choose **CDN Storage** . Otherwise, choose **Outside Limelight infrastructure**.

Table 1. Content Location Settings

> **Note:** If the **CDN Storage** option is unavailable, and you are already using CDN Storage, your service is not fully configured. If this is the case, please contact Limelight support.

Wizard Step: Basic Configuration

If you chose **CDN Storage** as your Content Location in Step 1, you'll see the following in the *Basic Configuration* step. Note that **Origin Hostname** is a drop-down menu that lists the available types of CDN Storage:

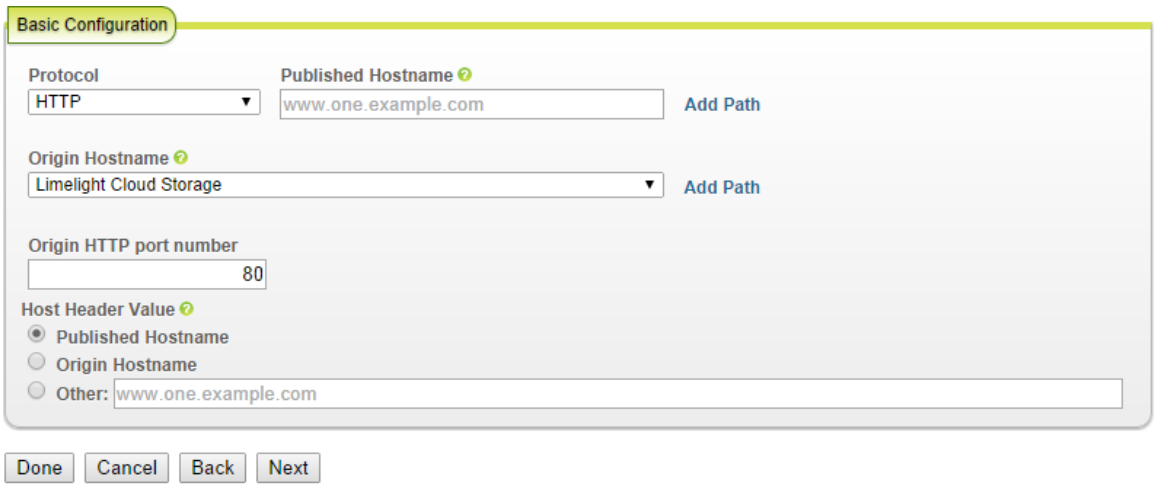

Figure 7. Basic Configuration Step - CDN Storage

If you chose **Outside Limelight infrastructure** as your Content Location in the *Content Location* step, you'll see a slightly different view. Note that **Origin Hostname** becomes an editable field, and you can also select an **Origin Protocol**:

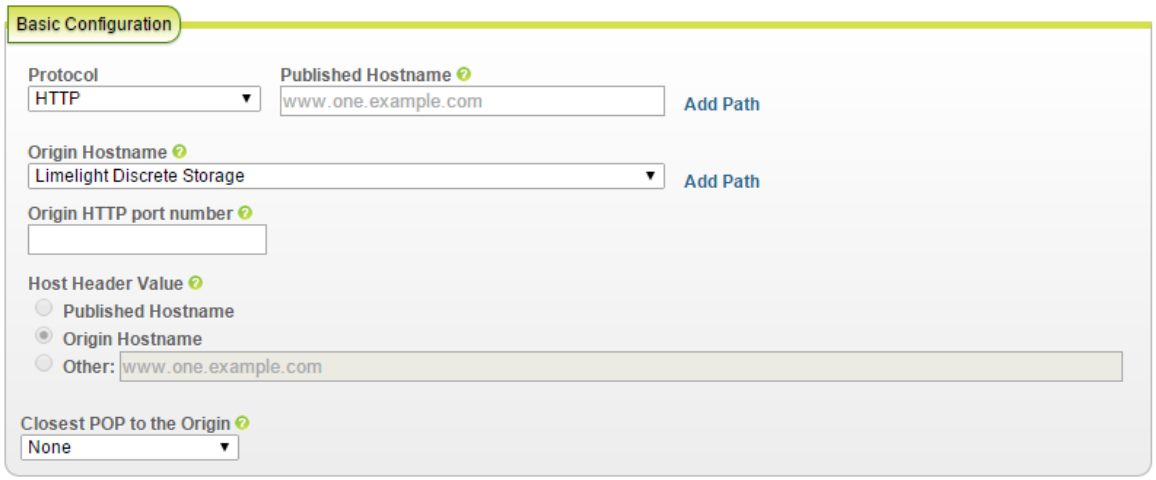

Figure 8. Basic Configuration Step - Outside Limelight Infrastructure

Protocol

The **Protocol** drop-down menu controls the level of security used when delivering content to your users.

Setting	Information Requested	Why It's Needed	Selecting the Right Option
Protocol	The HTTP protocol(s) to use to when delivering your cached content to end users	To ensure your content is delivered with the level of security you require	• To always deliver content insecurely, select **HTTP** • To always deliver content securely (via SSL), select **HTTPS** • To deliver content using the protocol specified in the incoming HTTP request, select **Both HTTP and**

			HTTPS

Table 2. Protocol Settings

Published Hostname

The **Published Hostname** is the domain name used in all public links (Published URLs) to your cached content.

Setting	Information Requested	Why It's Needed	Selecting the Right Option
Published Hostname	The fully-qualified domain name that will be used in all links to your cached content	To direct your users to the service (instead of your origin)	Enter the published hostname specified in your Welcome Letter, or a CNAME if desired, in the **Published Hostname** field . . Please note that IP addresses are not accepted. You must enter a fully-qualified domain name. If you want to use a directory name "alias" for a particular origin path, you can add the alias by clicking the **Add Path** link.

Table 3. Published Hostname Settings

Limelight will provide you with the correct **Published Hostname** in the *Welcome Letter* associated with your Limelight Account. By default, the **Published Hostname** will be in the form similar to: `accountname.vo.llnwd.net`

If you prefer to publish a different hostname name, you can use a DNS CNAME record to alias (point) your desired name to the one provided by Limelight.

Origin Protocol

The **Origin Protocol** drop-down menu controls how content is requested from your origin (when the content is not found in cache or has expired in cache).

Setting	Information Requested	Why It's Needed	Selecting the Right Option
Origin Protocol	The HTTP protocol(s) to use when retrieving content from your origin	To ensure your content is retrieved with the level of security you require	• To always retrieve content insecurely, select **HTTP** Always • To always retrieve content securely (via SSL), select **HTTPS Always** • To retrieve content using the protocol specified in the user's HTTP request, select **Match Inbound Protocol**

Table 4. Origin Protocol Settings

Origin Hostname

> **Note:** this option appears only when you select **Outside Limelight Infrastructure** as your content location.

Setting	Information Requested	Why It's Needed	Selecting the Right Option
Origin Hostname	The fully-qualified domain name or IP address of your origin server	The *Chunked Streaming* service needs to know where to get your content when users first request it, and also when it needs to be refreshed in the cache	Enter the domain name or IP address of your origin server in the **Origin Hostname** field. Please note that if you enter a domain name, it must be fully qualified. If your content is all in particular path on your origin, or you added a directory name "alias" with the Published Hostname for a particular origin path, you can enter the origin path by clicking the **Add Path** link.

Table 5. Origin Hostname Settings

Origin HTTP port number

Origin HTTP Port Number is the web server port Limelight will use in association with your Origin Hostname. The default for HTTP is port 80, and this value is pre-filled in the **Origin HTTP Port Number** field. The default for HTTPS is 443, and this is the value used by Limelight for all HTTP requests to origin (the value is not editable).

Setting	Information Requested	Why It's Needed	Selecting the Right Option
Origin HTTP port number	The HTTP port number to use when communicating with your origin server	If you are using a port other than the default (80) for HTTP, the *Chunked Streaming* service needs to know which port you've chosen	Leave the default port number for HTTP unless you are using another port number. If so, enter the new port number in the **Origin HTTP Port Number** field.

Table 6. Origin HTTP Port Number Settings

Host Header Value

Host Header Value specifies the value that the service will include in the HTTP `Host` header when making requests to your origin.

Setting	Information Requested	Why It's Needed	Selecting the Right Option
Host Header Value	The value to include in the HTTP `Host` header when communicating with your origin server	To help prevent end users from requesting content directly from your origin.	If you plan to block requests to your origin based on the value of the `Host` header, select **Published Hostname** or enter a value in the **Other** field If you chose **Limelight Storage** as your content location, the **Host Header Value** defaults to the **Origin Hostname**.

			Note: If you are hosting more than one origin on a single server, please see the additional information below.

Table 7. Host Header Settings

Closest POP to Origin

Setting	Information Requested	Why It's Needed	Selecting the Right Option
Closest POP to the Origin	Whether the *Chunked Streaming* service should always request origin content using a specified group of Limelight POPs	In some cases, performance can be improved by specifying POPs	This option is available only when configured by Limelight. Please consult Limelight Support if you are not sure which POP to choose. If you chose **Limelight Storage** as your content location, **Closest POP to the Origin** is preselected for best performance.

Table 8. Closest POP to Origin Settings

Browsers usually include the origin domain name of the requested URL in the HTTP `Host` header. You can use this behavior to detect and block such requests on your origin, denying those with a `Host` header that matches your domain name, and accepting those that match either your **Published Hostname** or another value you enter in the **Other** field.

If you are hosting more than one origin on a single server and you want to block based on `Host` headers, don't use **Published Hostname** - enter a value in the **Other** field instead. If you are hosting more than one origin on a single server and you don't want to block based on `Host` headers, choose **Origin Hostname**.

Example Settings

Configuration Field	Value	Notes
Protocol	HTTPS	Accept only HTTPS requests for cached content
Published Hostname	`published.host.com`	Use a CNAME alias instead of the name provided in the Welcome Letter (need to set up the CNAME separately)
Add Path	`/pubimages/`	Use the pubimages directory to uniquely identify the content in cache
Origin Protocol	HTTP Always	Always use HTTP to communicate with the origin server
Origin Hostname	`origin.host.com`	
Add Path	`/images/`	Directory path to the origin content; note that this doesn't need to match the path (if any) for the Published Hostname
Origin HTTP port number	`80`	Use the default HTTP port (no need to change anything)

Host Header Value	Published Hostname	This will block most browser requests made directly to origin

Table 9. Example Settings

Using the example configuration settings above, if `favicon.ico` is not cached for this configuration, or has expired in cache, a request to `https://published.host.com/pubimages/favicon.ico` will result in an origin request for `http://origin.host.com/images/favicon.ico`, with an HTTP `Host` header of `published.host.com`.

Object	Incoming Request	Origin Request

favicon.ico	`https://published.host.com/pubimages/favicon.ico`	`http://origin.host.com/images/favicon.ico`
		
		`Host: published.host.com`

Table 10. Example Response

Wizard Step: Basic Cache

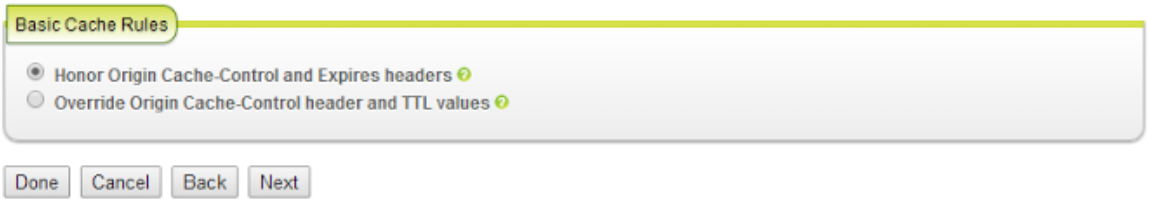

Figure 9. Basic Cache Step - Honor Origin Settings

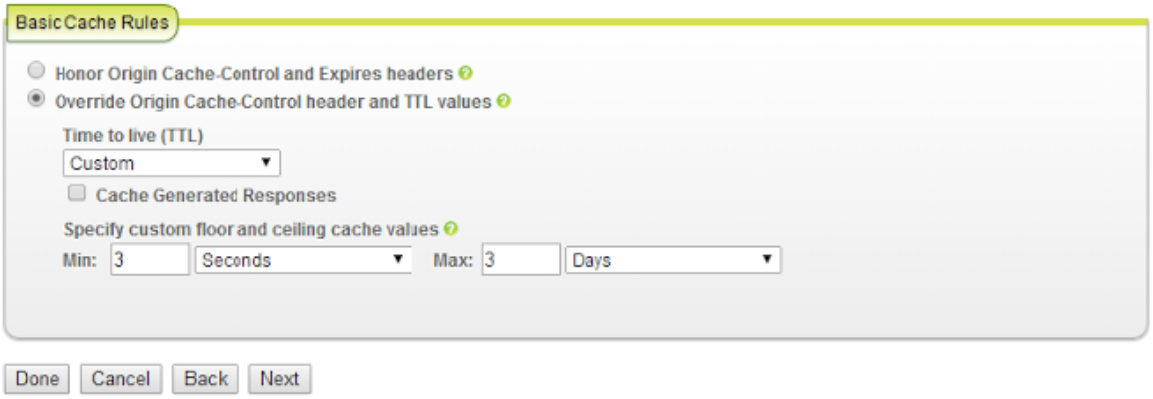

Figure 10. Basic Cache Step - Override Origin Settings

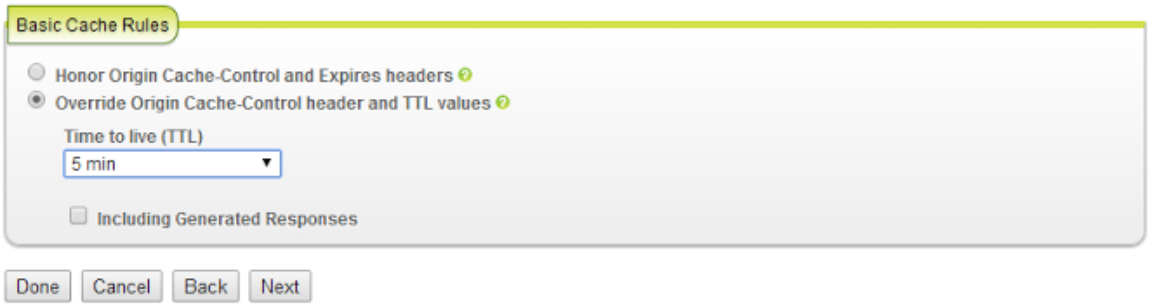

Figure 11. Basic Cache Step - Override Origin Settings - Custom TTL

Setting	Information Requested	Why It's Needed	Selecting the Right Option
Basic Cache Rules	Whether to override the default method for determining if an object in cache is expired	In some cases you may want to take explicit control over object expiration times (TTL - "Time To Live").	To allow Chunked Streaming to calculate TTL, select **Honor Origin Cache-Control and Expires headers** . Otherwise, choose **Override Origin Cache-Control header and TTL values** . If you want to set TTL to a specific length of time, select one of the times in the Time To Live (TTL) drop-down menu. Otherwise, to allow adaptive TTL calculation, select Custom from the Time To Live (TTL) drop-down menu. **Note:** TTL values set in this step are applied to all content under the **Published URL**, except for chunks and manifest files - TTLs for these are set individually in the *Media Delivery* step. In particular, TTL values set in this step apply to Cross Domain (`crossdomain.xml`) and Client Access Policy (`clientaccesspolicy.xml`) files, which should be in the root directory of the Published URL.

Table 11. Basic Cache Settings

Use the **Honor Origin Cache-Control and Expires headers** setting unless you have a specific reason to override the way in which object Time To Live (TTL) is calculated.

By default, *Chunked Streaming* considers an object "stale" (expired from cache) if the number of seconds specified by the associated `Cache-Control: s-maxage` or `Cache-Control: max-age` header has elapsed since initial caching or since the last freshness check, or if neither header is present, if the date and time in the `Expires` header has passed. The order of precedence is `Cache-Control: s-maxage` , `Cache-Control: maxage` , then `Expires`.

If no explicit freshness information is supplied (there are no `Cache-Control: s-maxage`, `Cache-Control: max-age` or `Expires` headers), and a `Last-Modified` header is present, the CDN will by default use the adaptive cache freshness algorithm to calculate remaining TTL, based on 20% of the age of the cached response, subject to a floor of 3 seconds and a ceiling of 3 days.

If you need to override (ignore) the above behavior, you can use the **Override Origin Cache-Control header and TTL values** option to specify a new TTL value using the **Time to live (TTL)** drop-down menu.

You can also control whether generated responses are cached using **the Cache Generated Responses** checkbox (for the **Custom** option) or **Including Generated Responses** (for other values in the drop-down menu).

Note:
Generated responses are HTTP responses that are generated dynamically ("dynamic content"). These responses often do not include any of the cache control headers needed to determine TTL, and are not cached by default to avoid caching personalized or user-specific responses.

By default, Limelight defines a generated response as one that is missing all of the following headers:

- `Expires`
- `Last-Modified`
- `Cache-Control: max-age`
- `Cache-Control: s-max-age`

If you choose the **Custom** option for **Time to live (TTL)**, you can change the parameters of the cache freshness algorithm using **Specify custom floor and ceiling cache values**.

If desired, the floor (minimum) can be raised and the ceiling (maximum) can be lowered or raised. If min and max are set equal to each other, the TTL becomes explicit, rather than adaptive.

Wizard Step: Media Delivery

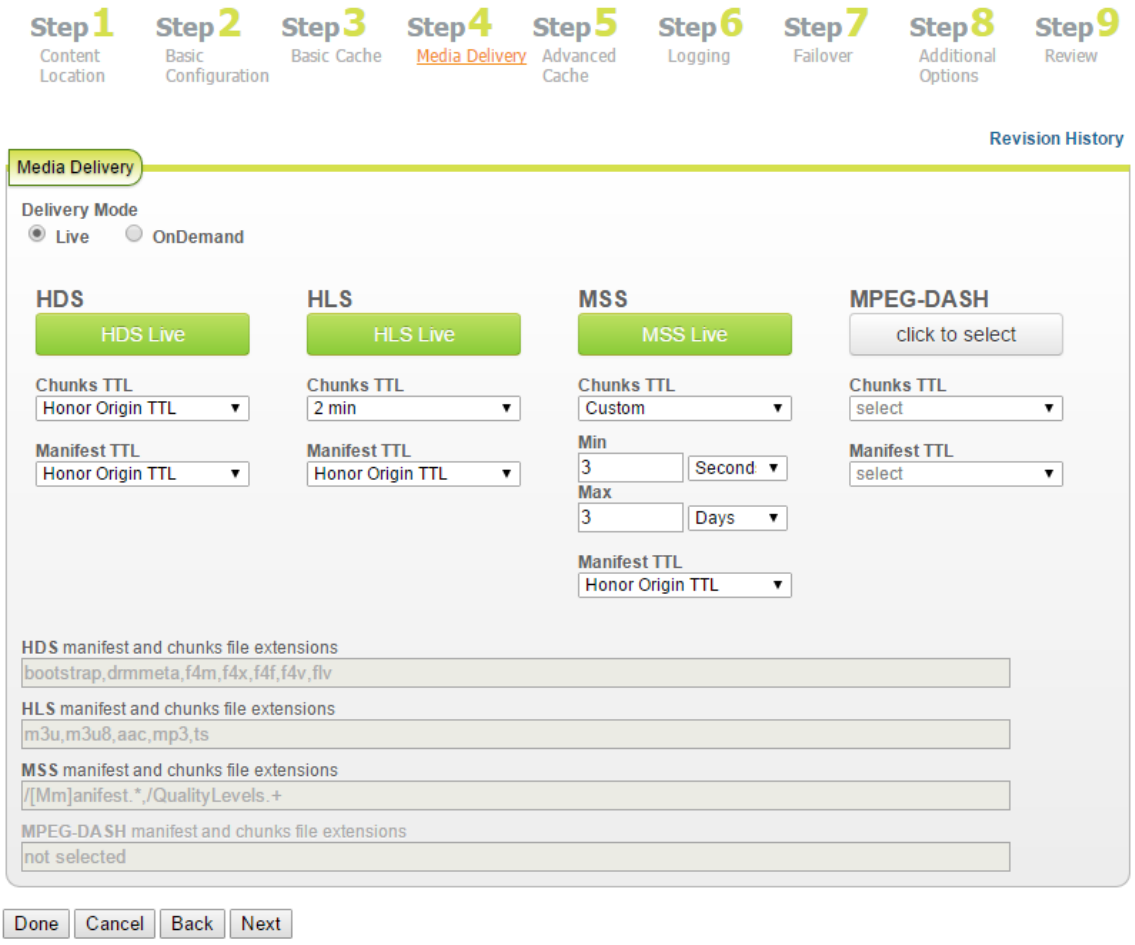

Figure 12. Media Delivery Step

Setting	Information Requested	Why It's Needed	Selecting the Right Option
Delivery Mode	Whether your configuration is for live video streams, or for video files available at any time (on demand)	To properly configure manifest file cacheability	Choose **Live** if you are delivering live streams, or **On Demand** for files that can be accessed at any time.
HDS / HLS / MSS / MPEG-DASH	The video formats to use when delivering your content	To properly configure optimizations for each format	Select each format (**HDS** , **HLS** , **MSS** or **MPEG-DASH**) that you want to deliver, based on the requirements of your video client.
Chunks TTL	The amount of time each chunk will be cached	In some cases you may want to take explicit control over object expiration times (TTL - "Time To	If the TTL provided by your origin is correct, use the default **Honor Origin TTL** setting. Otherwise, choose a specific TTL value (2 minutes to 60 days), or select **Custom** to set both the minimum and maximum TTL.

		Live")	
Manifest TTL	The amount of time the manifest will be cached	In some cases you may want to take explicit control over object expiration times (TTL - "Time To Live")	If the TTL provided by your origin is correct, use the default **Honor Origin TTL** setting. Otherwise, choose a specific TTL value (2 minutes to 60 days), or select **Custom** to set both the minimum and maximum TTL.
manifest and chunks file extensions	The file extensions you used for manifest files and content chunk files, for each media type you selected	The *Chunked Streaming* service needs this information to function properly	Please use the file naming rules below for manifest and chunk files. These formats must be followed for the *Chunked Streaming* service to work properly: **HDS** — **Chunks**: use Seg/Frag annotation in the file names; **Manifest**: use one of the following file extensions: `bootstrap`, `drmmeta`, `f4m`, `f4x` **HLS** — **Chunks**: use one of the following file extensions: `aac`, `mp3`, `ts`; **Manifest**: use one of the following file extensions: `m3u`, `m3u8` **MSS** — **Chunks**: include the keyword `QualityLevels` in the URL; **Manifest**: include either `Manifest` or `manifest` as a keyword in the URL **MPEG-DASH** — **Chunks**: use the `dash` file extension; **Manifest**: use the `mpd` file extension

Table 12. Media Delivery Settings

Wizard Step: Advanced Cache

Advanced Cache

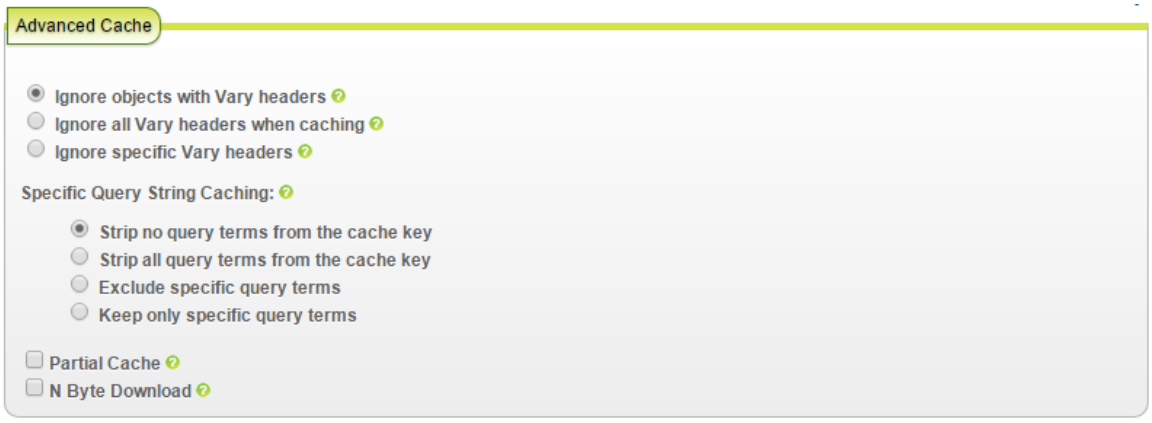

Figure 13. Advanced Cache Step - Advanced Cache Controls

Setting	Information Requested	Why It's Needed	Selecting the Right Option
Vary Headers	Which `Vary` response header fields *Chunked Streaming* should use when differentiating versions of an object in cache	*Chunked Streaming* stores a separate version of a requested object for each unique set of request header fields specified by the `Vary` header. If the `Vary` header specifies request header fields that change frequently, multiple copies of the same object may be stored in cache. To control this behavior, you can configure *Chunked Streaming* to ignore all `Vary` headers or specific `Vary` headers when caching and retrieving objects. All of the Vary headers associated with the object are still maintained and passed on to the client in the response.	• If you only want to cache a single version of an object regardless of its `Vary` header fields, choose **Ignore objects with Vary headers** • If you want to cache a new version of an object whenever any of its `Vary` header fields changes, choose **Ignore all Vary headers when caching** • If want cache a new version of an object whenever all but certain specified `Vary` header fields change, choose **Ignore specific vary headers** and select the `Vary` headers fields to ignore

Table 13. Advanced Cache Settings

Optimization

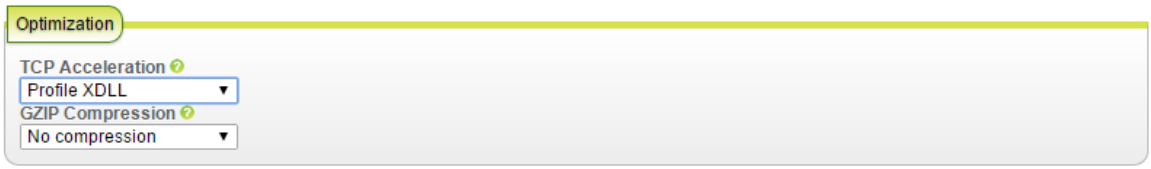

Figure 14. Advanced Cache Step - Optimization Controls

Setting	Information Requested	Why It's Needed	Selecting the Right Option
TCP Acceleration	The "profile" to use when accelerating the transfer of IP packets by modifying default TCP parameters	In certain circumstances, you may want to change the **TCP Acceleration** profile to optimize your delivery performance	When **TCP Acceleration** is enabled, the XDLL profile is the most efficient in many cases. **Note: TCP Acceleration** is an advanced configuration setting, and should only be changed if you're an expert user.
Gzip Compression	Whether to use Gzip compression when delivering XHTML, JavaScript, CSS, and other text files	Compressed objects are delivered more quickly, potentially improving the user experience	• If you want to provide all compressed files from your origin server, choose the **Gzip Passthrough** option. • If you prefer to have the *Chunked Streaming* service compress files when the requesting client can accept them, choose **Gzip on-the-fly** . • If you need to modify Gzip compression defaults, choose **Custom** , then either **Gzip on-the-fly** or **Gzip Passthrough** , and enter your Gzip modification extensions • You can also choose **No compression** if none of your files should be delivered compressed. For more information on this feature, see Gzip Details. **Note:** Compression cannot be applied to chunks.

Table 14. Optimization Settings

Gzip Details

When **Gzip Passthrough** is enabled, and a client indicates (via HTTP request header) that it prefers to receive compressed content, the *Chunked Streaming* service will serve a compressed version of the requested object if one is available on the origin server.

> **Note:** ***Gzip Passthrough*** *is available to all customers. If it is not enabled for you, please contact Limelight Support.*

If **Gzip On-the-fly** is selected, the *Chunked Streaming* service creates, caches, and delivers Gzip-compressed content as needed.

Compressible file types include: `action, ashx, asmx, asp, aspx, axd, cfm, css, css3, csv, do, doc, docx, htm, html, js, jsf, json, jsp, php, portal, rtf, svg, svgz, tsv, txt, xhtml, xml, site root (/),` and extensionless URLs.

Request and Response Headers

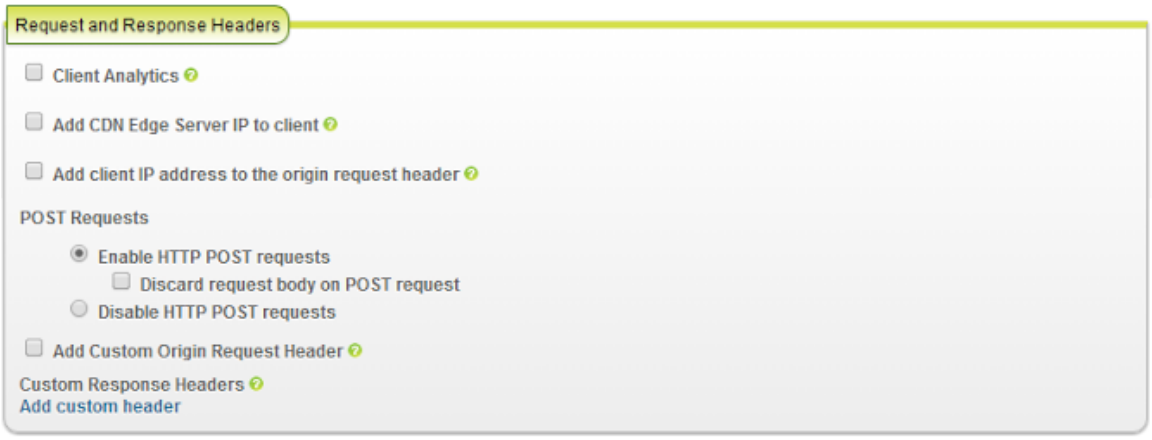

Figure 15. Advanced Cache Step - Header Controls

Setting	Information Requested	Why It's Needed	Selecting the Right Option
Client Analytics	Whether you want the *Chunked Streaming* service to provide geographic user information when requesting content from your origin	You may want to internally capture, analyze and report on user geographic information.	To use this feature, check the **Client Analytics** checkbox. The geo information is provided to your origin server via two request headers: `X-IP-Geo-Country` and `X-IP-Geo-All` . The geo fields provided are `continent`, `country` , `state`, `city`, `dma_id` and `asn`.
Add CDN Edge Server	Whether to provide clients with the IP	If you are using a custom client to display content, and you are also	To enable this feature, check the **Add CDN Edge Server IP to**

IP to client	address of the *Chunked Streaming* Edge Server responding to their requests	capturing performance-related data via the client, you may want to include the *Chunked Streaming* Edge Server IP address for later analysis and reporting. The IP address will be provided in the X-IP-Address response header.	**client** checkbox
Add client IP address to the origin request header	Whether you want the *Chunked Streaming* service to provide the requesting client's IP address in a custom header when requesting content from your origin	You may want to internally capture, analyze and report on user IP information	To enable this feature, check the **Add CDN Edge Server IP to client** checkbox, and enter the header name(s) that should contain the client IP address. The default header name is `True-Client-IP` . Note that the above headers are in addition to `X-Forwarded-For` , which is always provided to the origin.
POST Requests	Whether you want to accept or ignore POST requests from clients	If you are using a custom client to display content, you may want to allow it to communicate analytics or other information to your origin. Alternatively, you may want to convert POST requests to GET requests, or simply ignore them.	• To ignore all POST requests, select **Disable HTTP POST requests** . *Chunked Streaming* will respond with an HTTP 413 "Request Entity Too Large" status code to all POST requests. • To accept POST requests and pass them through to your origin, select **Enable HTTP POST requests** . If a POST request body exceeds 500 MB, *Chunked Streaming* will respond with an HTTP 413 "Request Entity Too Large" status code. • To accept POST requests but treat them as GET requests, select **Enable HTTP POST requests,** and check the **Discard request body on POST request** checkbox. POST bodies will be discarded.
Custom Request Headers	Whether you want to include custom headers and values whenever *Chunked Streaming* makes a request to your origin	If you want to tag all requests from *Chunked Streaming* for later analysis	To add a custom origin request header, click the **Add custom header** link, and enter a unique header name and value

Custom Response Headers	Whether you want to include custom headers and values whenever *Chunked Streaming* responds to a client request	If you are using a custom client to display content, you may want to provide it with information that uniquely identifies the *Chunked Streaming* service, Limelight Account, etc.	To add a custom client response header, click the **Add custom header** link, and enter a unique header name and value

Table 15. Header Settings

Progressive Video Download

Note : Progressive Video settings are displayed only when necessary for compatibility with legacy configurations.

Figure 16. Advanced Cache Step - Progressive Video Download Controls

Setting	Information Requested	Why It's Needed	Selecting the Right Option
Enable seek in FLV files	Whether to allow video client to skip forward and back (seek) within FLV files based on parameters specified in the query terms of the request URL.	Custom clients may want to provide the "seek" capability ("forward" and "back" buttons)	To enable this feature, check the Enable seek in FLV files checkbox
Enable H264 seek in MP4 files	Whether to allow video client to skip forward and back (seek) within MP4 files based on parameters specified in the query terms of the request URL.	Custom clients may want to provide the "seek" capability ("forward" and "back" buttons)	To enable this feature, check the **Enable H264 seek in MP4 files** checkbox

Table 16. Progressive Video Download Settings

Content Security

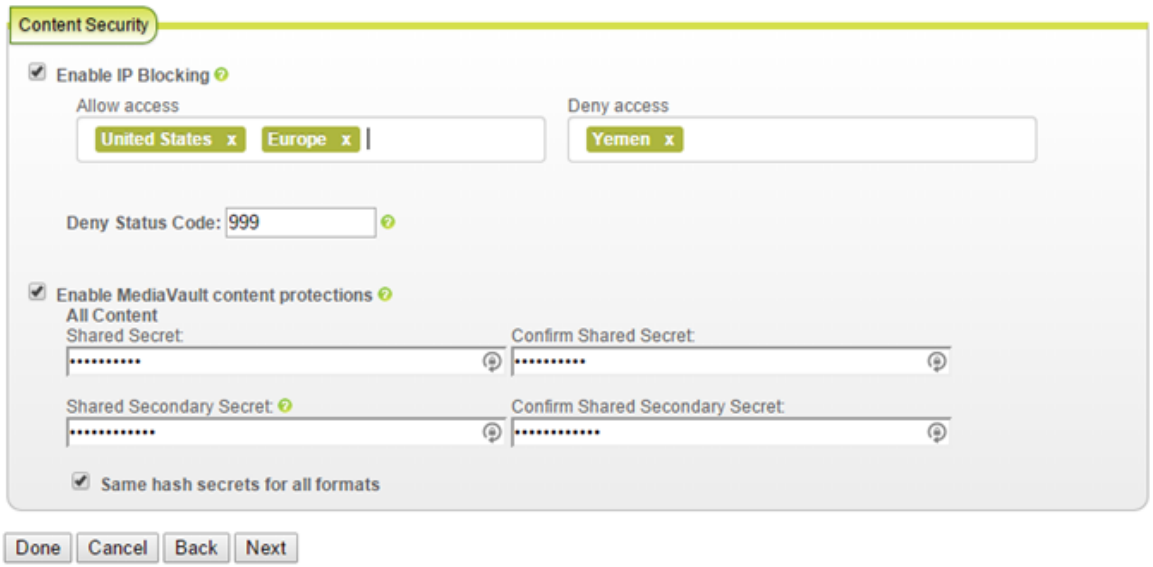

Figure 17. Advanced Cache Step - Content Security Controls

Setting	Information Requested	Why It's Needed	Selecting the Right Option
Enable IP Blocking	Whether you want to allow or deny access to your content by geographic region or named IP list	You may need to limit the audience for your content to certain regions to meet licensing restrictions, or exclude certain regions that are outside your target audience.	Once you have determined which geographies you want to manage, begin typing the name of a continent or country in the desired field (**Allow access** or **Deny access**). As you type, a pop-up list will display available choices that match your entry. If named IP lists have been created for your Account, you can also enter those names in the fields. If you need to provide the video client with an error code when access is blocked, enter an integer code in the **Deny Status Code** field.
Enable MediaVault content protection	Whether you want to use to provide additional content security. provides high-performance cookie-based URL authentication.	*MediaVault* can help you prevent "deep linking" and other unauthorized viewing behavior	To enable this feature, check the **Enable MediaVault** content protection checkbox, and provide a primary and secondary "shared secret" (both used to prevent URL tampering). If you need to provide different shared secrets for each media format, uncheck the **Same hash secrets for all formats** checkbox. Separate fields will then be

			displayed for each media format. For more information, see Details

Table 17. Content Security Settings

MediaVault Details

MediaVault is a high-performance URL authentication service. *MediaVault*'s main purpose is to help you secure your content from unauthorized viewing.

MediaVault maximizes authentication performance by using tokens to avoid three-way handshakes (common to other methods of authentication) that can lead to severe connection time latency.

Please note that *MediaVault* is *not* a replacement for DRM and should not be associated with user authentication.

MediaVault works like this:

- You enter a shared secret during the configuration process
- You then generate a token (MD5 hash) for each published URL, based on the shared secret, and append it to the URL in a query term or provide it in a cookie.You can generate the token manually by navigating to the *Configure* > *MediaVault* in the navigation pane, or by creating server-side code on your origin.
- When a request is received, *MediaVault* uses the same hash algorithm to create it's own token, which should be identical to the one you appended.
- If the tokens match, *MediaVault* then looks for additional *MediaVault* -specific query terms (such as end date/time and IP address/mask) to determine whether the request is valid. If the tokens don't match, the URL was tampered with and the request is rejected.

For more information, see the MediaVault User Guide by navigating to Help Center > Documentation > Delivery > Guides > MediaVault in the navigation pane.

> **Note:** *MediaVault* does not support IPv6-based Client IP or IP Range restrictions at this time.

Wizard Step: Logging

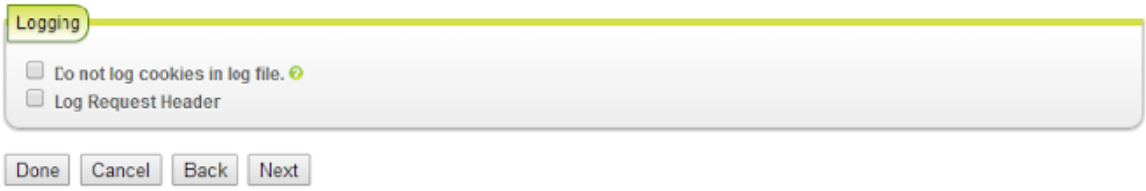

Figure 18. Logging Step Controls

Setting	Information Requested	Why It's Needed	Selecting the Right Option
Do not log cookies in log file	Whether you want *Chunked Streaming* to *stop* saving cookie information in your log files	If you process log files and don't need the information in the Cookie header, you may want to remove it to simplify processing and/or reduce log file size.	If you know you *don't need* Cookie header information in your log files, check the **Do not log cookies in log file** checkbox. Otherwise, leave it unchecked.

Log Request Header	Whether you want *Chunked Streaming* to *start* saving specific Request Headers in your log files	If you process log files and need access to information in the Request Headers, you may want to enable this option	If you know you *do need* Request Header information in your log files, check the **Log Request Header** checkbox and enter the names of the specific headers to log. Otherwise, leave it unchecked.

Table 18. Logging Settings

Chunked Streaming normally logs all `Cookie` header information, up to a maximum of 8 KB for the entire header (regardless of the number of cookies in the header).

Wizard Step: Failover

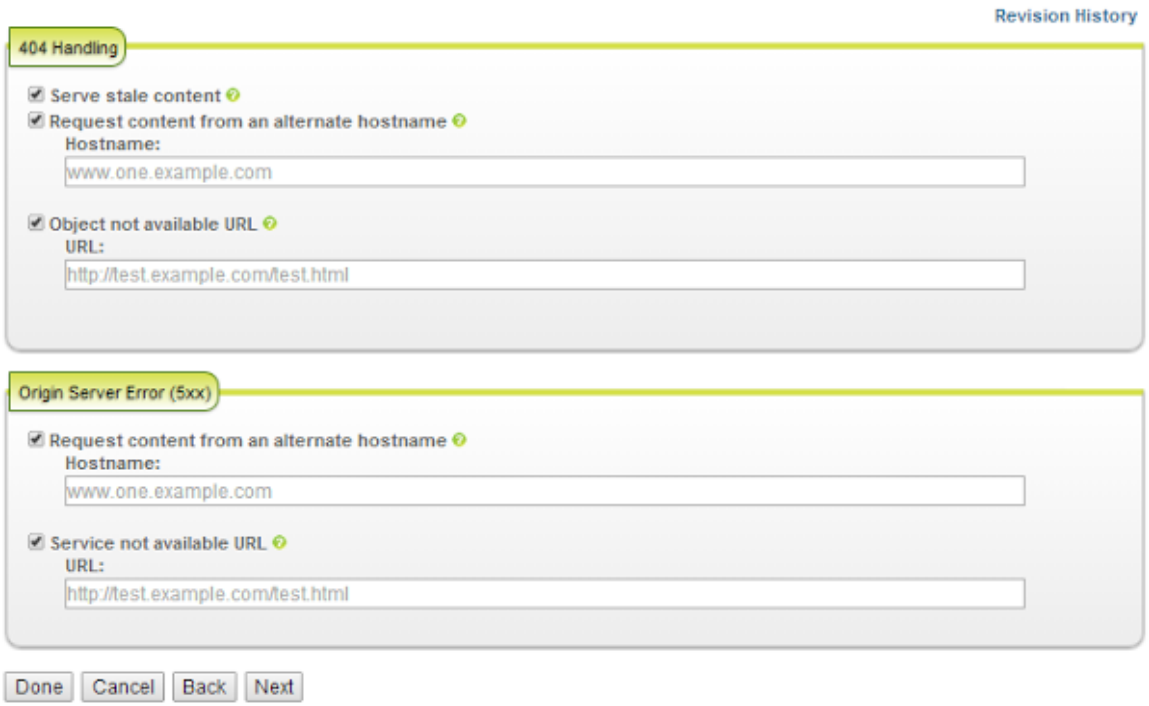

Figure 19. Failover Step Controls

Normally, when the CDN receives a `404` response from the origin, it is passed back to the requesting client. With the **Serve stale content** option, if there is cached content for a given request and the origin returns a `404`, the Edge Server returns the stale content instead of issuing a `404` to the client.

Object not available URL specifies a URL to be used for content retrieval instead of sending a `404` response code to the client. The original request is reissued to the fallback URL with any modifications still in place. The option is used to provide a custom "not found" message or to provide instructions to retry the request. Any URL used for this option must match its own configuration within the CDN.

Service not available URL specifies a URL to return instead of sending a `503` (Service Unavailable) or `504` (Gateway Timeout) response. The `503` or `504` may have been generated by the origin, but could also be internally generated by CDN should a connection failure to origin occur. Any URL used for this option must match its own configuration within the CDN.

Setting	Information Requested	Why It's Needed	Selecting the Right Option
Serve Stale Content	If the requested content is cached but stale (expired), and there is an HTTP `404` status when requesting a fresh version from your origin, whether you want to pass the `404` status back to the client, or serve the stale content instead	If an object has expired in cache, and your origin server returns a `404 (Page Not Found)` error when *Chunked Streaming* attempts to get a fresh copy of the object, you may want to serve the expired object instead of allowing the client to handle the `404` message.	If it's not acceptable for the client to handle the `404` message, and you are OK with serving stale content instead, check the **Serve Stale Content** checkbox. Note that if there is no cached object, a `404` message will still be returned to the browser.
Request content from an alternate hostname	If there is an HTTP `404` status when requesting fresh content from your origin, whether to try a backup origin before handling the `404` status	If your primary origin returns a `404` status, and you have a backup origin, you may want *Chunked Streaming* to try the backup before handling the error	To try a backup origin if the primary origin responds with a `404` status, enable **Request content from an alternate hostname** and enter the fully-qualified hostname of the backup origin. Note that specific ports are not supported.
Object not available URL	Whether you want to pass HTTP `404` status messages back to the client, or serve a custom error page instead	If an object has expired in cache, and your origin server returns a `404` error to *Chunked Streaming* , you may want to serve a custom error page instead of allowing the client to handle the 404 message.	If you want to take control over the content displayed by clients when there is a `404` from origin, check the **Object not available URL** checkbox, and enter the fully-qualified URL of the content to serve.
Request content from an alternate hostname	If there is an HTTP `5xx` status when requesting fresh content from your origin, whether to try a backup origin before handling the `5xx` status	If your primary origin returns a `5xx` status, and you have a backup origin, you may want *Chunked Streaming* to try the backup before handling the error	To try a backup origin if the primary origin responds with a `5xx` status, enable **Request content from an alternate hostname** and enter the fully-qualified hostname of the backup origin. Note that specific ports are not supported.
Service not available URL	Whether you want to pass HTTP `503` and `504` status messages back to the client, or serve a custom error page instead	If an object has expired in cache, and your origin server returns a `503 (Service Unavailable)` or `504 (Gateway Timeout)` error to Chunked Streaming , you may want to serve a custom error page instead of allowing the client to handle the 5xx message.	If you want to take control over the content displayed by clients when there is a `503` or `504` error from origin, check the **Service not available URL** checkbox, and enter the fully-qualified URL of the content to serve.

Table 19. Failover Settings

Wizard Step: Additional Options

Figure 20. Additional Options Step (Read Only)

You can use the **Additional Options** step to view any advanced configuration changes Limelight makes to your configuration.

If one or more such configurations is changed from its default value by Limelight, the **Additional Options** tab becomes visible, and the advanced configurations and their settings are displayed:

Setting	Information Requested	Why It's Needed	Selecting the Right Option
(various)	(none) This is a read-only display of advanced *Chunked Streaming* configuration changes Limelight has made to your configuration	The information in the Additional Options step can help you better understand your configuration.	If you have questions about any settings in Additional Options, please contact your Account Manager or Limelight support.

Table 20. Additional Options Settings (Read Only)

The advanced configuration options which can be configured for you by Limelight (and become visible in the *Additional Options* step) include:

Option Name	Description
Assume cacheable pending origin response	If an origin request is pending for an object, continue serving the object from cache
Cache entire object if range request less than offset	Cache the entire object for Range requests ending before the specified Byte offset
Cache hit/miss response trigger	Returns HIT or MISS in the X-CDN-Cache response header when the specified request header (trigger) is present
Cache only "popular" objects	Cache only objects that are "popular" based on the specified "points" (the approximate frequency an object is requested, in seconds)
Convert URL ranges to Range requests	Convert URLs ending in /range/x-y or /range/x- to origin GET range requests

Deny requests with specified Referrer header(s)	Deny requests with the specified Referrer header(s)
Disable object caching	Do not cache objects
Disable persistent origin connections	Disable persistent origin connections ("enabled" is the default global configuration)
Do not add max-age on all requests to origin	Don't add Cache-Control: max-age=259200 header on origin requests (but do include any existing Cache-Control headers)
Enable partial caching by regex	Enable partial caching for object URLs that match the specified regex
Gzip compression level	Set the Gzip compression level (0 to 9). The default (and recommended) level is 1.
Ignore bad status codes from origin	Ignore bad status codes from origin (40x and 5xx). If FALSE, other rewrite options may redirect the client to specific URLs based on the status code.
Lowest allowed rate-limiting bitrate	Set the lowest bitrate allowed when rate limiting, in KBytes/second
Make cached URLs case-insensitive	Make the URLs of cached objects case-insensitive by converting all characters to lowercase in the Cache Key. When using this feature, all Purge requests must not contain any uppercase characters.
Max duration client can be idle while receiving response	After this time passes, the client is disconnected and the request is aborted. The default is 30 minutes.
Maximum object TTL	Set the maximum TTL value for all cached objects, in seconds, but honor Cache-Control headers if present
Minimum object TTL	Set the minimum TTL value for all cached objects, in seconds, but honor Cache-Control headers if present
Object TTL for "negative" origin response	Set the object TTL, in seconds, when there is a negative origin response (status codes other than 200, 203, 300, 301 and 401 and/or Cache-Control or Pragma headers with certain values). This rewrite overrides other origin cache control headers.
Origin connect timeout duration	Set the timeout, in seconds, for initiating origin connections (how long to wait when trying to establish a connection)
Origin reply timeout duration	Set the timeout, in seconds, for origin replies (how long to wait for a reply from origin)
Persistent client connection duration	Set the duration, in seconds, of persistent client connections
Persistent origin connection duration	Set the duration, in seconds, of persistent origin connections
Redirect clients to source URL	Redirect clients to the source URL with the specified status code
Refresh-check cached content on every request	Check for fresh origin objects (newer versions of objects) on every request. Most commonly used in conjunction with *Ignore bad status codes from origin* to enable the origin to allow or deny every request by inspecting all request parameters, including Cookies.
Remove specified response header(s)	Remove origin response headers that match the specified value
Retry failed *MediaVault* HTTPS hash	If an HTTPS *MediaVault* hash check fails, retry the same hash-check URL using HTTP

checks	
Store *MediaVault* hash in cookie	Keep the *MediaVault*hash secret in a browser cookie (rather than in a URL parameter)
Treat empty responses with 200 status as 404 status	Treat “empty” origin responses (no content body) with 200 status codes as if they are 404 status codes

Table 21. Additional Options Available

Wizard Step: Review

The Review step is the final step in the configuration process, and gives you an opportunity to confirm you’ve made the changes you intended before submitting your configuration.

Setting	Information Requested	Why It’s Needed	Selecting the Right Option
(all changes from previous steps)	Review and approve the changes you made in previous steps	To confirm you’ve changed the settings you expected	Click **Submit** to save your settings in a new configuration. Click **Back** to make changes to previous steps, or **Cancel** to discard all of your unsaved settings.
Notes	Optional notes that you can refer to later when browsing historical configuration changes	You may find it helpful to include additional information for others (why the configuration changes were made, etc.)	If you want to save notes with your configuration, just enter them in the Notes field
Revision History	Optional. Limited availability.	You can review all of your previous configuration changes, and “roll back” to any previous state if desired	This feature is accessible via the "Revision History" link in any wizard step

Table 22. Review Fields

Revision History Details

If you select the "Revision History" link in any wizard step, you can now review all of your previous configuration changes, and roll back to any previous state.

The configuration version number, creation date/time, submitting user, and any notes entered in the Review step are displayed for each historical change

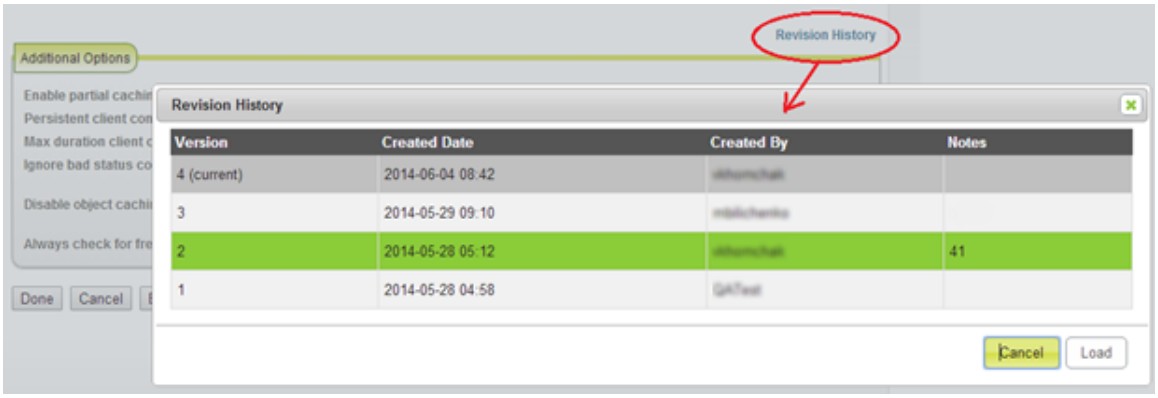

Figure 21. Revision History link and pop-up dialog with historical configuration versions

Selecting a version loads its settings into the configuration wizard and displays an alert above the wizard tabs to make it clear that a historical configuration is loaded.

As with any configuration changes, clicking **Submit** in the *Review* step will save the configuration, and clicking Cancel will discard the changes.

Editing a Chunked Streaming Configuration

To **Edit** a configuration:

- In the main menu, select **Configure**, select **Chunked Streaming**, choose a configuration to edit from the *Chunked Streaming Configurations* list, and then click **Edit**. The first step of the configuration wizard will appear.
- Configure each step as necessary then click **Submit**. For more information on the contents of each step, see Creating a New Configuration.

Copying a Chunked Streaming Configuration

To **Copy** a configuration:

- In the main menu, select **Configure**, select **Chunked Streaming**, choose a configuration from the *Chunked Streaming Configurations* list, and then click **Copy**. The first step of the configuration wizard will appear.
- Configure each step as necessary then click **Submit**. For more information on the contents of each step, see Creating a New Configuration.

> **Note:** All configuration options are copied, including those visible only in the *Additional Options* step. In the simplest case, you may only need to change the **Published Hostname** field in the *Basic Configuration* step.

Deleting a Chunked Streaming Configuration

To **Delete** a configuration:

- In the main menu, select **Configure**, select **Chunked Streaming**, choose a configuration from the *Chunked Streaming Configurations* list, and then click **Delete**. The *Please confirm the delete request* dialog appears.
- Type DELETE in the space provided to permanently delete the selected record and the configuration is deleted.

Viewing a Chunked Streaming Configuration

To **View** a configuration:

- In the main menu, select **Configure**, select **Chunked Streaming**, choose a configuration from the *Chunked Streaming Configurations* list, and then click **View**.

Configuring Chunked Streaming (v2)

Using the Limelight Control portal, you can manage a configuration that controls several manifests and Chunked Streaming origins. Chunked Streaming provides a way for you to adopt Limelight's optimized configuration profiles for delivering chunked video content through the CDN over HTTP, HTTPS, or both. To use Chunked Streaming, you first need to chunk your content and generate the associated manifest files (Chunked Streaming does not encode, transcode or transmux your media). You can host your content on your origin servers or with Origin Storage.

> **Note:** In general, Chunked Streaming is like Caching & Delivery, but Chunked Streaming allows you to create multiple delivery configurations for media formats.

Chunked Streaming List Page

Navigate to **Configure** > **Chunked Streaming (v2)** in the navigation pane. The *Configurations for* page is displayed, which contains a list of your Chunked Streaming configurations.

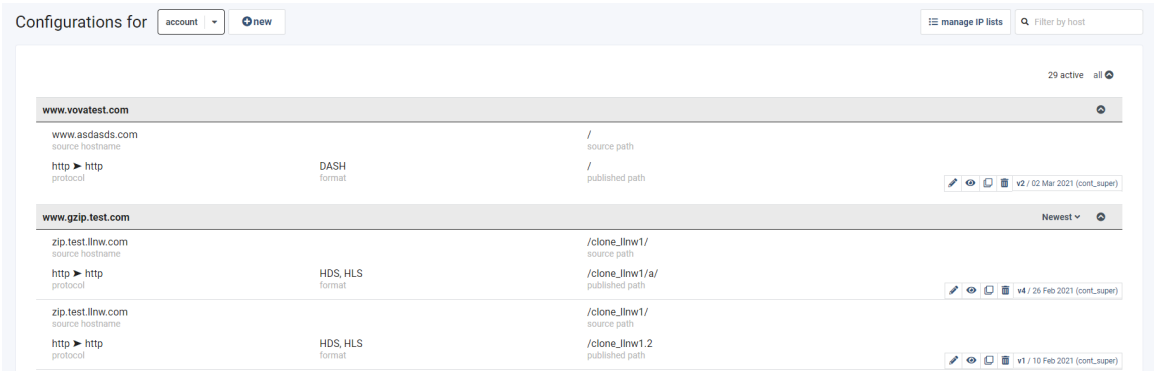

Figure 22. Configurations Page

The top gray bar shows the number of active configurations on the right. You can click the bar to hide or show all configurations.

The rest of the gray bars show individual configurations grouped by Published Hostname in the bar header. The Published Hostname is the public URL prefix used in links to your published content (URLs seen by end-users)

The *Chunked Streaming Configurations* list provides the following information for each configuration:

- **source hostname** (Origin Hostname) – The private URL prefix used by Limelight to retrieve and cache content from your origin server (not visible to end-users).
- **source path** – The private URL prefix used by Limelight to retrieve and cache content from your origin server (not
- visible to end-users).
- **protocol** – protocols in this format:
  `{source protocol} ➤ {published protocol}`
  Multiple protocol combinations are allowed. An example is: `HTTP ➤ HTTPS & HTTPS ➤ HTTP`
- **format** - The media formats which the content is delivered in
- **published path** - The path portion of the URL (visible to end-users).
- **action icons** - edit, view, clone, delete
- **configuration information** - version, date version created or last modified, and the user that created or modified the version

Filtering the List of Configurations

Use the **Filter by** drop-down menu and the filter text field to filter the list by specific fields:

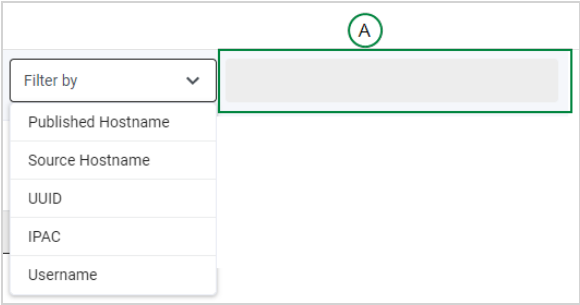

A - filter text field, initially disabled

To filter the list:

1. Make a selection in the drop-down menu.
2. Enter a value in the filter text field.
3. Press the Enter key on your keyboard.
   The list is reduced to include only configurations that match the filter.

Display the original list by clicking the **x** icon in the filter text field:

Creating a New Configuration

A Chunked Streaming configuration consists of one root instance and two child instances (manifest and chunk) per each media delivery format (HDS, HLS, MSS, DASH). All root and child instances can also support multiple combinations of protocols.

> **Note:** The root configuration serves as the anchor for all child video format configurations under it. The root specifies the base path from which all child configuration paths will be relative. The root specifies a base set of options set on each child configuration unless otherwise overridden by format-specific settings.

Service Profiles

Each new configuration is based on a Service Profile. Service Profiles define the configuration structure and specify default and mandatory options that must be applied on every configuration. A Service Profile can serve as both a guide and a guardrail for the type of content your configuration will serve (characterized by a Use Case).

The **Use Case** and **Service Profile** drop-down menus are disabled:

- In existing configurations.
- After you have selected a Published and Source Protocol while you are creating a new configuration.

> **Note:** If you have not already saved the new configuration but you want to choose another Service Profile, you can do so by exiting out of the **Create configuration** screen and creating a new configuration by clicking the **+new** button.

If you wish to modify a Service Profile or migrate, add, or remove a Protocol Set for an existing configuration, contact your Account Manager.

Creating a New Configuration

1. Begin the new configuration by clicking the **+ new** button at the top left of the page.
2. Make configurations in the sections in the subsequent *Create configuration* screen.
3. When you are finished, click the Activate button at the bottom of the page or the floating blue checkmark icon on the right side of the screen.

Links to instructions for sections on the screen:

Content Location

Media Delivery

Chunked Streaming Options Applicable to Both Root and Child Instances

Others

Additional Options

Notes

Content Location

With two exceptions, configuring a Chunked Streaming *Content Location* is similar to Caching & Delivery (v2).

- When configuring the Source URL path, keep in mind that it will be the root for all child video format configurations. For example, you have a DASH configuration, and you configure `/streaming/` as the source URL path, then that path will have two children:
  `/streaming/DASH/manifest`
  `/streaming/DASH/chunk`
- **The path ends with a filename** and **Only publish files with these extensions** options are not present under the Source URL path field in the Chunked Streaming *Content Location* section.

For additional information, see Content Location in the Caching & Delivery (v2) section.

Media Delivery

In some cases, you might want to take control over cached object expiration times. Chunked Streaming Delivery allows you to indicate how long (TTL) the manifest and chunks will be cached.

You can make configurations for any supported Chunked Streaming formats (HDS, HLS, MSS, DASH) in this UI section.

Adding a Media Format

1. Click the **Add format** button and choose a format from the subsequent drop-down menu.
   A configuration section for the selected format is displayed.

   > **Note:** Several additional sections are also displayed below the *Media Delivery* section. (Instructions for the additional UI sections are described later in this document.)

2. Configure the Chunk TTL and Manifest TTL fields.
   Both fields allow you to configure manually, default values to the root instance, or choose from preset times.

Option	Description/Instructions
Default to Root	The TTL defaults to the root instance value, configured in the Caching Rules section immediately below this one.
Configure manually	1. Configure values in the **Min** and **Max** fields. 2. Configure a TTL for cached responses in the Cache Generated Responses field.
preset value	The value you select will be the expiration time.

Cache Generated Responses Field

For HTTP Chunked Streaming responses generated dynamically, origins often do not supply cache-control headers (`Cache-Control`, `Expires`, or `Last-Modified`). This field directs EdgePrism to consider such responses cacheable. Configure a number and time unit (seconds, minutes, etc.) to keep the responses in the cache.

Removing a Media Format

- If you are working with a new configuration, hover your mouse pointer over the right side of the section and click the **x** button that appears. The configuration is removed, and the removal cannot be undone.
- If you are editing an existing configuration, click the **x** button to remove a format. If you click the **< back** link at the top, respond 'No' to the prompt asking if you want to discard changes.

Chunked Streaming Options Applicable to Both Root and Child Instances

Although the following sections of the user interface apply to root and child Chunked Streaming instances, the fields and how you interact with them are identical to Caching & Delivery. Please refer to Caching & Delivery documentation as indicated in the following table.

Section in the Chunked Streaming User Interface	Caching & Delivery Documentation Reference
Caching Rules	Caching Rules
Arc Light	Arc Light
Optimization	Optimization
Headers & Methods	Headers & Methods
Failover	Failover
Content Security	Content Security
Logging	Logging
Cookie Handling	Cookie Handling
Redirect	Redirect

For information about other sections not in the preceding table, see:

Others

Additional Options

Others

This section presents additional delivery options you can use in the Chunked Streaming configuration. For descriptions, hover your mouse pointer over the right side of the option name. An information icon appears along with a description of the option.

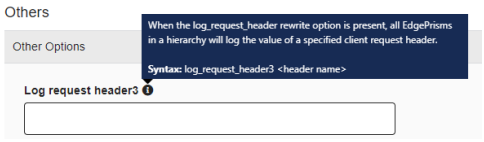

Figure 23. Additional Options Description

Additional Options Section

This section in the UI allows you to fine-tune options for the formats you selected in the Media Delivery section. By default, several significant options are present that you can modify or delete. You can add additional options by type all or part of an option name in the **Type option name** field, where partial matches are considered.

Each option is paired with the protocol sets you selected during earlier configurations. You can configure an option at the root level or any (and all) children and children in each format you configured. Each Child has child entries for manifest and chunk.

To configure an option:

1. Click the menu icon in an option's row.
2. Select a combination of root and children in the subsequent popup menu.
   To show the children click the show link at the bottom of the popup menu.

To delete an option, hover your mouse pointer over the right side of the section and click the **x** button that appears. The configuration is removed, and the removal cannot be undone.

Notes Section

Enter any notes related to the configuration.

Editing a Configuration

On the *Configurations for* page, select the **edit** icon on the right side of the row.

> **Note:**
> On rare occasions, a configuration might contain unsupported protocol set configurations, and if you attempt to edit the configuration, Control prevents you from editing and displays this message:
>
> *"This protocol combination is not supported in this application. You may view it here, but to make modifications, either do so via our configuration API or reach out to your account team."*
>
> Unsupported protocol sets are generally the byproduct of migrating a configuration from an older configuration version.

Viewing a Configuration

On the *Configurations for* page, select the desired row or select the **preview** icon on the right side of the row.

Cloning a Configuration

On the *Configurations for* page, select the **clone** icon on the right side of the row.

> **Notes:**
> When you clone a configuration, you must, at a minimum, enter a new Published URL Path.
>
> A cloned configuration requires about 15 minutes to be published. Until then, it is in the 'Pending' state.

Deleting a Configuration

If you just created a new configuration:

1. On the *Configurations for* page, select the **delete** icon on the right side of the row.
2. Confirm your intention to delete in the subsequent dialog.

Reverting to a Previous Configuration

Each time you update a configuration, a new version is assigned.

To revert to a previous configuration:

1. On the *Configurations for* page, select the **revert** icon on the right of the row.
   A list of previous versions is displayed in a dialog.
2. Select the version to which you want to revert.

   > **Note:** Although you intend to revert to a previous version, the reverted version will become the current version with a new version number. The new version number is displayed at the bottom of the dialog.

3. Click the **Activate** button.

Configuring DNS Services

About DNS Services

DNS Services provide an easy-to-use, DNS-based, global load balancer used for directing end-user requests to customer Resources, for example, web servers.

DNS Services are managed by the Director, which is a DNS service that can route traffic based on IP address, end-user nameserver geographic location or the BGP autonomous system number of the end-user nameserver.

Failovers, also known as 'Traffic Balancers', redirect traffic in case a Resource is not available.

To access the DNS Services page, navigate to **Configure** > **DNS Services** in the navigation pane.

DNS Services Entities

DNS Services comprises the entities in the following table.

Entity	Description
Director	The Director is a DNS service that helps balance and manage end-user requests to origin servers and other IP Resources (including requests from more than one CDN). For example, if you have end users in diverse geographic locations, the Director provides content specific to their region with the best site performance and end-user experience. The Director can also block traffic from a country, province, or IP address. The Director can route traffic based on: • IP address • end-user nameserver geographic location • BGP autonomous system number (ASN) of the end-user nameserver.
Resources	Resources are IP addresses or hostnames that you want to manage. Resources include zero, one, or at most two Health Checks. Resources optionally participate in Failovers.
Failovers	If a Resource is inaccessible due to a network failure or errors in Resource configurations, you can define a Failover for the Resource in which one or more other Resources act as the "Failover" Resources. Each Resource in a Failover has a relative priority and preference; both are positive integers. **Note:** • A value of '1' is the highest preference and subsequently greater numbers indicate decreasing preference. • Priority, used for unequal load balancing, works in reverse. Larger values mean the record will get served more frequently relative to the other records with the same preference but

Entity	Description
	smaller weight. Failovers have two modes of operation. • In **single Failover** mode, records with the lowest preference values are replaced with the next available record in order of preference. (A value of '1' is the highest preference and subsequently greater numbers indicate decreasing preference.) A lower preference record will only be used if there are not enough higher preference records available to satisfy minimum response (see Adding a New Failover). If there are no available records, failed records will continue to be deactivated until minimum response is reached. If all records are in a failure state, records from the highest preference set are returned until inactive records become available. • In **group Failover** mode, the active set of records is treated as a single unit. Records are deactivated on failure and not replaced with lower preference records. When the active record set is less than the minimum response, the entire active record set is deactivated. The next preference set is activated if the number of active records in that set is greater than or equal to the minimum response. This process will continue until no preference set is available for activation. The last active set will remain active until a preference set has enough available records to meet the minimum response threshold. The settings determine how Failovers monitor and react to failure conditions. Failovers include zero, one, or at most two Health Checks. Failovers include at least one Resource.
Director Policies	A Director Policy assigns a Time to Live (TTL) and weight to an alias host and canonical host, and optionally binds the Policy to a business Rule called a Director Policy Rule.
Director Policy Rules and Match Components	Director Policy Rules have a name and description and at least one Match Component that determines the mode for routing a request: by country, region, ASN, or CIDR.
Health Checks	You can add Health Checks to Failovers and Resources. • Failovers Part of the Failover configuration process is adding Health Checks. Health checks let you easily add or remove a Resource from all of your Policies without going through each Policy. This may be helpful if you are adding new Resources into the rotation and do not want to activate them right away or in a case where you want to remove a Resource out of the handout rotation (due to maintenance or other factors). Health Checks also detect unavailability quickly before your users are seriously impacted.

Entity	Description
	• Resources Sometimes Resources become unavailable and Health Checks can detect this quickly before unavailability seriously impacts your users.

Page Layout

The page provides tabs for working with three components:

- **Resources** tab - the IP addresses or hostnames that you want to manage.
- **Failovers** tab - Failovers in case a domain is not available
- **Director policies** tab - binding one or more Resources together using a business Rule

Configuration Overview

Three core steps need to be performed to configure DNS Services:

- Adding and Configuring Resources: During this step, you will add, configure, and indicate to Limelight which Resource(s) you wish to manage. Ultimately, the Resources are the “handout answers” to end-user query requests for the hostname. See Working with Resources.
- Adding and configuring Failovers: During this step, you will configure the Resources in a Failover group. See Working with Failovers.
- Adding and configuring Policies: During this step, you will add and configure one or more Policies. The Policy is the act of binding one or more Resources together using a business Rule for the distribution of end-user requests to your added and configured Resources. See Working with Director Policies.

Changes made to DNS Services will propagate to the CDN edge in less than 10 minutes; however, the Freshness value may dictate how quickly changes are ultimately acquired, and traffic will begin to shift.

Working with Resources

To work with Resources, click the **Resources** tab. The *Resources* list displays:

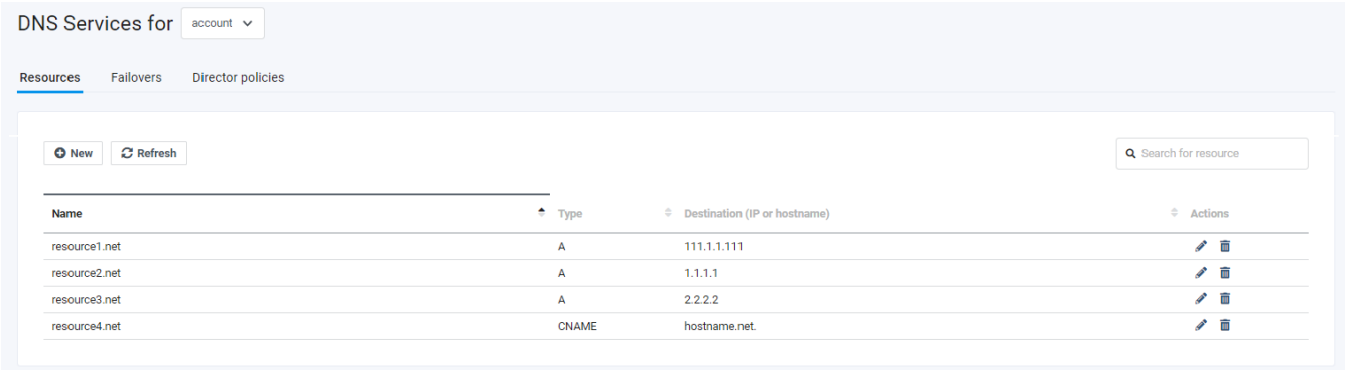

The *Resources* list contains the columns in the following table for each Resource.

Column	Description
Name of Resource	Name that was given to the DNS Resource when it was added

Column	Description
Type	Record category - IP (A record) or host name (CNAME)
Destination (IP or hostname)	IP (A record) or host name (CNAME) that the Resource points to.
Actions	Icons for editing and deleting a Resource.

Searching for DNS Services Resources

- Begin typing search criteria in the **Search for resource** field at the top right of the list.
  As you type, matches are highlighted in yellow and only rows containing matching characters are displayed.
- To view all rows, remove the search criteria from the **Search for resource** field.

Note:
You can search for Resources by the following columns:
- **Name**
- **Destination (IP or hostname)**

Searches are case-insensitive.

Adding a New DNS Services Resource

The first step in configuring DNS Services is to add one or more Resources to which you want to direct traffic.

To add a new Resource**:**

1. Select the **Resources** tab. The *Resources* list displays.
2. Select the **+ New** button on the right side of the page under the tabs header, and the *Create resource for* page displays
3. Complete fields on the page (see Fields on the 'Create resource for' Page).
4. Click **Save**
   You are returned to the Resources tab and a message is displayed stating that a job to create the Resource has started.
   Click the **Refresh** button periodically to determine if the Rule has finished processing.

Notes:
- After configuring a Resource you may have to wait five or more minutes for the Resource to be created.
- The Resource will be added to the list only after the job is complete.
- After creating the Resource, you might want to add Health Checks to it. See Adding Health Checks to a DNS Services Resource.

Fields on the 'Create resource for' Page

Field	Description/Instructions
Name	Unique alphanumeric name for the new Resource. If you are adding a large number of Resources, you may want to establish an easy-to-understand naming scheme to make it easier to find specific Resources in the *Resources* list.

Field	Description/Instructions
Type	Resource category: A - DNS address record. Maps a domain name to an IP address. CNAME - DNS Canonical name record. Maps a domain name to a Canonical name record.
Target (IP address or Domain name)	The IP address or fully-qualified domain name of the Resource. Enter the IP address (only IPv4 addresses are allowed) or fully-qualified domain name of the Destination (IP or hostname) Resource. If you enter a hostname instead of an IP address, the query response to any balanced or directed group will be returned as a CNAME, not an A Record. **Notes:** • If the domain name does not include a terminating dot ("."), the zone name is appended to the domain name when the resource is pushed to the edge. For example, if the zone name is `lldns.net.` and the domain name is specified as `www2.lldns.net`, the name is `www2.lldns.net.lldns.net.` when pushed to the edge. • To add hostname `www3` to zone `lldns.net.` you can specify the domain name as either `www3.lldns.net.` or `www3` Both naming methods result in the record being pushed to the edge as `www3.lldns.net.`

Adding Health Checks to a DNS Services Resource

To add a Health Check:

1. Locate the desired Resource on the **Resources** tab, then select the pencil (edit) icon.
2. Click the **+ Add new healthcheck** button.

   Note: If the Resource already has both types of Health Checks, the **+ Add new healthcheck** button is disabled.

   The *CREATE HEALTHCHECK* dialog displays.
3. Complete the fields in the dialog (see Fields In the 'CREATE HEALTHCHECK' Page).
4. Click the **Save** button.
   The dialog closes and a message is displayed stating that a job to create the Health Check has started.
   Click the **Refresh** button periodically to determine if the Rule has finished processing.

   Notes:
   - After configuring a Health Check you may have to wait five or more minutes for the Resource to be created.
   - The health will be added to the *Edit resource for* page only after the job is complete.

Fields in the 'CREATE HEALTHCHECK' Dialog

Field	Description/Instructions
Type	Health Check category that you want the Failover to react to. If you've already added one type, only the other type is available to select. Select the type: • **ping**: Select this value to check for a response from the Resource itself. • **http**: Select this value to check for a response from an object served by the Resource.
Check interval (seconds)	Time span between Health Checks, in seconds. Defaults to 60 if not specified. Enter a value greater than or equal to 60.
URL	For http Health Checks only, the relative URL of the object that is served by the Resource.

Editing a DNS Services Resource

To edit a Resource:

1. Choose a Resource from the list in the '**Resources**' tab and then click the pencil (edit) icon in the *Actions* column.
2. Change the settings as required and click **Save**.
   A message is displayed stating that a job to save the Resource has started, and the word "processing" is displayed on the right side of the row. Click the **Refresh** button above the list periodically to determine if the Rule has finished processing.

For more information about the fields, see Adding a New DNS Services Resource and Adding Health Checks to a DNS Services Resource.

Deleting a DNS Services Resource

To delete a Resource:

1. Choose a Resource from the list in the '**Resources**' tab, and then click the trash can (delete) icon.
2. In the confirmation dialog, click **OK**.
   A message is displayed stating that a job to delete the Resource has started, and the word "processing" is displayed on the right side of the row. Click the **Refresh** button above the list periodically to determine if the Rule has finished processing.

Notes:
- If you accidentally delete a Resource, you must recreate it by adding a new Resource.
- When you delete a Resource, it is removed from all Failovers it participates in.

Working with Failovers

Searching for Failovers

- Begin typing search criteria in the **Search for failover** field at the top right of the list.
  As you type, matches are highlighted in yellow and only rows containing matching characters are displayed.
- To view all rows, remove the search criteria from the **Search for failover** field.

> **Note:** You can search for Resources by the **Name** column. Searches are case-insensitive.

Adding a New Failover

To add a new Failover:

1. Select the '**Failovers**' tab.
2. From the '**Zone**' drop-down menu, click the zone to which the Failover is to pertain.
3. Select the **+ New** button on the right side of the page under the tabs header, and the *Create failover for* page displays.
4. Complete fields on the page (see Fields on the 'Create failover for' Page).
5. Click **Save**.
   A message is displayed stating that a job to create the Failover has started, and the word "processing" is displayed on the right side of the row. Click the **Refresh** button above the list periodically to determine if the Failover has finished processing.

> **Note:** If you are at all unsure of the settings, do not attempt to change them yourself. Instead, contact your Account Manager.

Fields on the 'Create failover for' Page

Field	Description/Instructions
Name	Host name for the Failover Resource. If you enter a name without a terminating dot, the zone name will be appended to the host name before the Failover is sent to the edge. **Notes:** • If the domain name does not include a terminating dot ("."), the zone name is appended to the domain name when the resource is pushed to the edge. For example, if the zone name is `lldns.net.` and the domain name is specified as `www2.lldns.net`, the name is `www2.lldns.net.lldns.net.` when pushed to the edge. • To add hostname `www3` to zone `lldns.net.` you can specify the domain name as either `www3.lldns.net.` or `www3` Both naming methods result in the record being pushed to the edge as `www3.lldns.net.`

Field	Description/Instructions
Type	Category of Failover: 'Single' or 'Group'.
Minimum response	Minimum number of records the Failover group is to respond with.
Maximum response	Maximum number of records the Failover group is to respond with.
Delay time	Delay time in seconds to use before marking a Resource as 'in service'.
Percent Threshold	Threshold percentage of Health Checks that should pass to prevent triggering a Failover condition for a record. Enter a decimal number between 0.0 and 100.00
Health checks	Select the Health Check type that you want the Failover to react to: • To check for a response from the Resource itself, select '**ping**'. • To check for a response from an object that is served by the Resource, select **http**. **Note:** You can select either option or both options.
TTL	Time to live for the served record.
Resources	Resources that you designate as the Failover Resources. See Configuring Resources in a Failover Group.

Configuring Resources in a Failover Group

Any Resources that have been configured for the Failover group are listed in the '**Resources**' section of the **Create failover** or **Edit failover for** pages.

You can configure a Resource's priority and preference and add or remove a Resource.

Configuring Relative Priority and Preference

In addition to name, type, and target (see Fields on the 'Create resource for' Page), each Resource listed has a priority and preference:

- preference - The integer preference value for the Resource in the Failover group. A value of '1' is the highest preference and subsequently greater numbers indicate decreasing preference.
- priority - The integer priority value for the Resource in the Failover group. A positive integer to use for unequal load balancing. Larger integers mean the record will get served more frequently relative to the other records with the same preference but smaller weight.

Note: Resources are selected first based on preference, then priority.

To configure relative priority and preference:

1. Make entries in the **Priority** and **Preference** fields for each Resource:
   - Preference: Relative preference compared to other Resources.
   - Priority: Relative priority compared to other Resources.
2. Click **Save**.
   A message is displayed stating that a Failover job has started, and the word "processing" is displayed on the right side of the row. Click the **Refresh** button periodically to determine if the Failover has finished processing.

Removing a Resource

To remove a Resource, click the trash can (delete) icon on the right side of the Resource's row. The resource is removed from the **Resources** list.

You can also remove a Resource as described in Adding and Removing Resources from a Failover Group.

Adding and Removing Resources from a Failover Group

1. Click **+ Add/remove resources**.
   The RESOURCES dialog is displayed. Resources currently configured for the Failover group contain a minus (remove) icon. All other Resources contain a plus (add) icon.
2. Indicate the Resources you wish to add by clicking the **+** button on the right side of the Resource's row.
3. Add or remove Resources by clicking the plus or minus icon.
4. Click **Apply** at the bottom of the list to save your selections.
   The selections are added to or removed from the *Resources* section.

Editing a Failover

To edit a Failover:

1. Choose a Failover from the list in the '**Failovers**' tab, and then click the pencil (Edit) icon.
2. Change the settings as required (see Fields on the 'Create failover for' Page and Configuring Resources in a Failover Group).
3. Click **Save**.
   A message is displayed stating that a Failover job has started, and the word "processing" is displayed on the right side of the row. Click the **Refresh** button periodically to determine if the Failover has finished processing.

Deleting a Failover

To delete a Failover:

1. Locate the Failover in the list in the '**Failovers**' tab, and then click the trash can (delete) icon.
2. In the confirmation dialog, click **Delete**.
   A message is displayed stating that a job to delete the Failover job has started, and the word "processing" is displayed on the right side of the row. Click the **Refresh** button periodically to determine if the Failover has finished processing.

Working with Director Policies

You work with Director Policies in the **Director policies** tab.

The **Director policies** tab contains a list of configured Policies for a given zone. Each row in the list shows the information in the following table.

Field	Description
Alias Host Name	A DNAME that maps or renames an entire sub-tree of the DNS namespace to another domain.
Canonical Host	The CNAME of the host the DNAME points to.
Weight	Relative importance of this record compared to other records with the same Alias Host Name but different Canonical Host names.
TTL	Time to live. Amount of time the Policy will remain in the DNS resolver's cache.
Rule	Business Rule associated with the Policy. Click the name to edit the Rule. See Editing a Director Policy Rule.
Rule Description	Summary of the Rule.
Rule Match components	The actual matches on ASN, Country, and CIDR.
Action icons	Icons to edit and delete Policies. **Note:** If a Policy has been recently configured but has not finished processing, the word 'processing' shows instead of the 'edit' and 'delete' icons. Click the **Refresh** button periodically to determine of the Policy has finished processing.

Note:
Except for **Weight** and **TTL**, long field values are truncated. Do either of the following to see the complete description:

- Hover the mouse pointer over the field value.
- For all fields except **Rule**, double-click the field value, then copy and paste.

Searching for Director Policies

- Begin typing search criteria in the **Search for policy** field at the top right of the list.
  As you type, matches are highlighted in yellow and only rows containing matching characters are displayed.
- To view all rows, remove the search criteria from the **Search for policy** field.

Note:
You can search for Policies by the following columns:

- **Alias Host Name**
- **Canonical Host**
- **Rule**
- **Rule: Description**

Searches are case-insensitive.

Adding a New Director Policy

1. From the '**Zone**' drop-down menu, select the zone to which the Policy is to pertain.
2. Click the **+ New** button under the tab headers on the right side of the page.
   The *Create director Policy for* page appears.
3. Fill in the fields on the page (see Fields in the 'Create director Policy for' Page).
4. Click **Save**.
   A message is displayed stating that a Director Policy job has started, and the word "processing" is displayed on the right side of the row. Click the **Refresh** button periodically to determine if the Director Policy has finished processing.

> **Note:** After creating a new Policy, you must wait five or more minutes for the Rule to finish processing.

Fields in the 'Create director policy for' Page

Field	Description/Instructions
Alias Host Name	DNAME that maps or renames an entire sub-tree of the DNS namespace to another domain.
Canonical Host	CNAME of the host the DNAME points to. **Notes:** • If the domain name does not include a terminating dot ("."), the zone name is appended to the domain name when the resource is pushed to the edge. For example, if the zone name is `lldns.net.` and the domain name is specified as `www2.lldns.net`, the name is `www2.lldns.net.lldns.net.` when pushed to the edge. • To add hostname `www3` to zone `lldns.net.` you can specify the domain name as either `www3.lldns.net.` or `www3` Both naming methods result in the record being pushed to the edge as `www3.lldns.net.`
TTL	Time to live. Amount of time the Policy will remain in the DNS resolver's cache.
Weight	Relative importance of this record compared to other records with the same **Alias Host Name** but different **Canonical Host** names. Defaults to 0 if not specified. Larger integers mean the record will get served more frequently relative to the other records with the same name or smaller weight.
Rule	Business Rule associated with the Policy. See Applying Director Policy Rules for instructions. If

Field	Description/Instructions
	no Rule is specified, the record acts as the default response when there is no Rule match from other records.
Comments	Notes about the Policy.

Working with Director Policy Rules

Director Policy Rules determine the basis for routing a request.

Rule Type	Routing Type
Geo	Geographic-based. Traffic is routed based on the end-user nameserver geographic location: country or region (province/state).
ASN	Traffic is routed based on the BGP ASN (autonomous system number) of the end user's nameserver.
IP Address	Traffic is routed based on an IP address. If the end-user query matches the IP in the Policy, the Limelight nameserver platform routes based on the IP address.

Applying a Director Policy Rule

To apply a Rule:

1. While creating or editing a Director Policy, click the **+ Apply/Remove rule** button.
   The *RULES* dialog displays.
2. Click the + icon on the right side of the Rule's row.
3. Click **Apply**.
   The Rule is added above the **Apply/Remove rule** button in the *Create director policy for* page.

> **Note:** You can apply only one Rule per Director Policy.

Creating a New Director Policy Rule

To create a new Rule:

1. While creating or editing a Director Policy, click the **+ Apply/Remove rule** button.
   The *RULES* dialog displays.
2. Click the **Create rule** button at the top left of the *RULES* dialog.
   The *CREATE RULE* dialog displays.
3. Fill in the fields in the dialog. See Fields in the CREATE RULES and EDIT RULE Dialog for details.
4. Click **Apply** in the *RULES* dialog..
   The Rule is added above the **Apply/Remove rule** in the *Create director policy for* page.

Editing a Director Policy Rule

1. While creating or editing a Director Policy, click the **+ Apply/Remove rule** button.
   The *RULES* dialog displays.
2. Click the **edit** (pencil) icon in the desired Rule's row.
3. Fill in the fields in the dialog. See Fields in the CREATE RULES and EDIT RULE Dialog for details.

> **Note:** To remove a Match Component, click the **(-)** icon beside the Component's row.

4. Click **Apply** to save the changed Rule.
   Within the *RULES* dialog, a message is displayed stating that a job to create or update the Rule has started. Unlike other processing jobs, you won't see the word 'processing' until you click the **Refresh** button.
5. After the first time you click the button, click it periodically to determine if the Rule has finished processing.

> **Note:** If you attempt to modify the Rule before it is finished processing, a message is displayed warning you that you must wait until the Rule finishes processing.

Deleting a Director Policy Rule

To delete a Rule:

1. Click the **delete** (trash can) icon.
2. A message is displayed stating that a job to delete the Rule has started and the word "processing" is displayed on the right side of the row. Click the **Refresh** button periodically to determine if the Rule has been deleted.

Fields in the CREATE RULES and EDIT RULE Dialogs

Field	Description/Instructions
Name	Descriptive name for the Rule.
Enabled	You can create a Rule but not enable it for various reasons. If a Rule is not enabled, it will not be applied. Add a check mark if you want to enable the Rule.
Description	Rule explanation.
Match components	Rules have multiple Components and each Component can consist of one or more of the following: country, region, ASN, and CIDR. The fields in a Component appear in increasing specificity and are 'AND'ed together. The most specific Rule Component that can be matched is applied. **Note:** You can add multiple Components using the dialog. All are associated with the Rule.

Field	Description/Instructions
	1. Configure the following fields. • Select 'Country' and 'Region' for geographic-based routing. **Notes:** ◦ The 'Region' field is supported for only a subset of all countries. ◦ You can select 'Country' without selecting 'Region' but you must select 'Country' to select 'Region'. • Enter the ASN (Autonomous System Number) to route based on Border Gateway Protocol (BPG). • Enter CIDR (Classless inter-domain routing) to route based on IP address. 2. Click the + button to the right of the Component to add the Component. 3. Click **Apply**. A message is displayed stating that a job to create or edit the Component has started. 4. Create additional Components as desired.

Editing a Director Policy

To edit a Director Policy:

1. Choose a Policy from the list in the '**Director policies**' tab, and then click the pencil (edit) icon.
2. Change the settings as required (see Adding a New Director Policy).
3. Click **Save**.
   A message is displayed stating that a Director Policy job has started, and the word "processing" is displayed on the right side of the row. Click the **Refresh** button periodically to determine if the Director Policy has finished processing.

Deleting a Director Policy

To delete a Director Policy:

1. Choose a Resource from the list in the '**Director policies**' tab, and then click the trash can (delete) icon.
2. In the confirmation dialog, click **Delete**.
   A message is displayed stating that a Director Policy delete job has started, and the word "processing" is displayed on the right side of the row. Click the **Refresh** button periodically to determine if the Director Policy has finished processing.

A message is displayed stating that a Director Policy delete job has started, and the word "processing" is displayed on the right side of the row. Click the **Refresh** button periodically to determine if the Director Policy has finished processing.

Configuring *MediaVault* Hash Generator

The MediaVault hash generator is a tool you can use to:

- Create signed URLs on an ad-hoc basis
- Learn how MediaVault works so you can implement hashes within your own applications

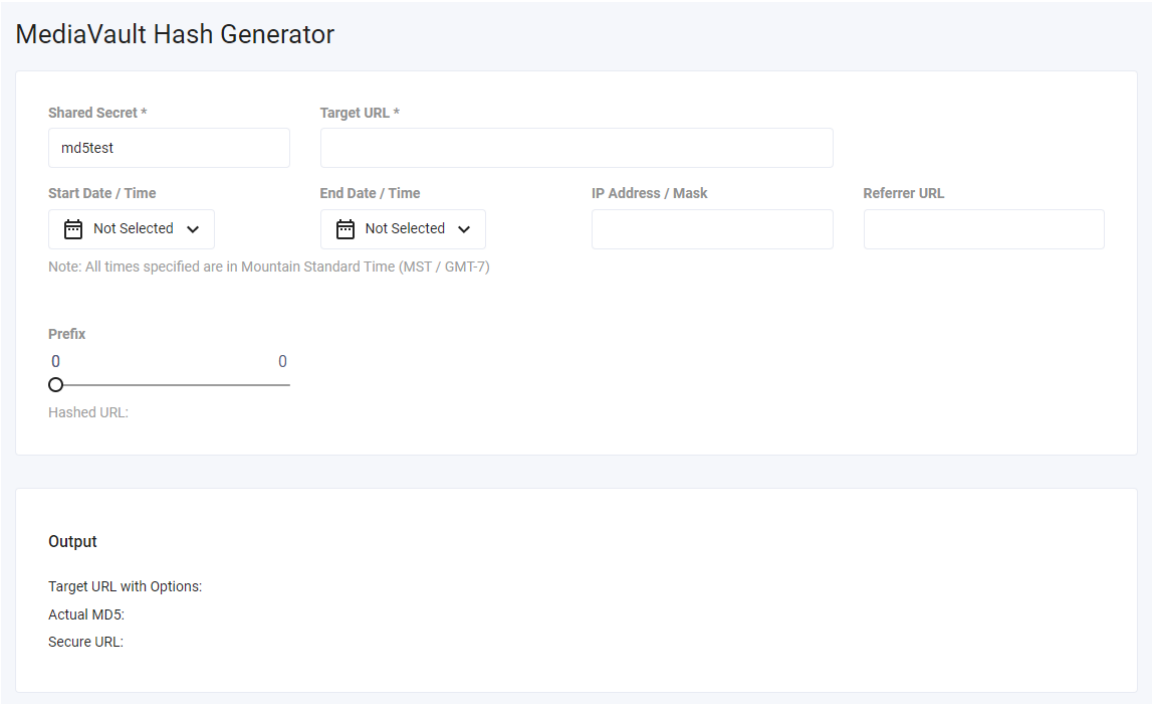

For additional information about MediaVault, and further configuration options, see the MediaVault Guide.

Instructions

1. For any given URL, fill out the fields in the top section of the page (See Configuration Fields). As you enter information, fields in the **Output** section are updated (see Output Fields).
2. Copy the value of the **Secure URL** field and publish it in place of the non-protected URL.

Configuration Fields

To gain greater insight into the purpose of all fields, see "MediaVault Parameters" in the MediaVault Guide.

Field	Description
Shared Secret	The shared secret that you have set on the respective CDN configuration. If you don't know your Shared Secret, refer to your CDN configurations in Limelight Control or contact Limelight Support.
Target URL	The URL you want to protect.
Start Date / Time	Optional

Field	Description
	The time the request is authorized from (represented as Unix epoch seconds). End users can't access the media before the start time Optional
End Date / Time	Optional The time the request is authorized to (represented as Unix seconds). End users can't access the media after the end time.
IP Address / Mask	Optional The IP address, IP address range, or Mask you wish to make your content available to.
Referrer URL	Optional Domain name from the Referer request header, usually the media player that is authorized to play your content.
Prefix	Optional For cookie-based MediaVault, the length of the Shared Secret + Target URL that should be hashed. See Using the Prefix Slider for instructions.

Using the Prefix Slider

The Prefix field has a slider that lets you generate a single hash to allow access to multiple objects nested under a single path. The slider value determines how much of the Shared Secret + Target URL should be hashed, and ranges from 0 to the length of the Target URL field.

Use the slider to generate a single hash to allow access to multiple objects nested under a single path.

Slider Value	How Much of Shared Secret + Target URL to Hash
0	Shared Secret + all of the Target URL **Note:** The URL that will be hashed is displayed as an example under the slider. This slider (option p) is added to omit hashing other options like IP address or referrer or just to strip the part of the URL.
1 through N	All of the Shared secret + N characters of the Target URL.

Example

Let **Shared Secret** = **md5test**

Let **Target URL** = **http://llnw.com/media-vault/media.mp4**

The slider maximum value is 37, the length of the **Target UR**L.

Slider Value	Result to Hash
0	md5testhttp://llnw.com/media-vault/media.mp4
1	md5testh
2	md5testht
37	md5testhttp://http://llnw.com/media-vault/media.mp4

Output Fields

Field	Description/Instructions
Target URL with Options	Informational only[1] The URL without hash specified, essentially the same as Secure URL but without `h` (hash) parameter.
Actual MD5	Informational only[1] Shows how the hash is generated, which is then passed to `h` parameter. It is formed as follows: `<md5secret><url><parameters>`.
Secure URL	The result of hashing: • Shared Secret • Target URL • Start Date / Time, End Date / Time, IP Address / Mas, and Referrer URL, if provided.

[1]These fields are provided in case you want to generate the secure URL yourself.

Configuring SSL Certificates

Certificate List Page

Navigate to Configure > SSL Certificates in the navigation pane. The SSL Certificates page is displayed, with a list of all certificates associated with the currently selected Limelight Account.

Action icons (view, edit, etc.) are displayed on the right side of each certificate's row.

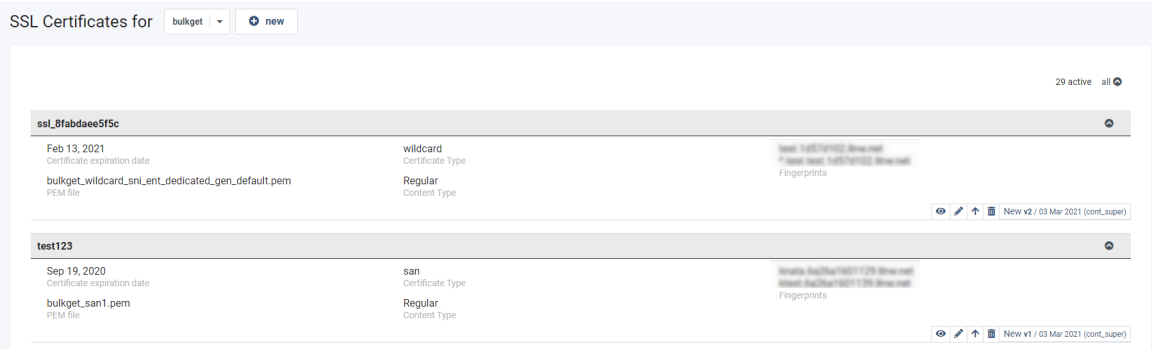

Note: Depending on a certificate's state, additional or fewer action icons may be present. All icons and capabilities are explained later in these instructions.

The page contains summary and detail information for the certificates listed.

Summary Information

The *SSL Certificates for* list provides this information:

Total of each type on the right side of the gray header bar. Possible types:

- new
- pending publication
- published
- updated
- expired
- withdrawn

Detail Information

Each certificate in the list provides this information:

Item	Description
Certificate name	Customer-assigned certificate name
Expiration Date	Certificate expiration date, extracted from the certificate's "Not After " element
Content type	Certificate purpose; possible values:

Item	Description
	• single server • SAN • wildcard
Fingerprints	List of domains covered in the certificate
PEM file	Certificate pem file name, if the certificate is PEM-encoded
Action controls	On right side of row; allow you to view, edit, publish, withdraw, and delete a certificate
Certificate state and version, date, user	To the right of the action controls on the right side of each row. Indicates the certificate version, date the certificate was created or updated, and user that created/updated the certificate

Working with SSL Certificates

You can create, view, edit, publish, withdraw, and delete your own Server Name Indication (SNI) SSL certificates in the Limelight Network.

Creating a New Configuration

Viewing Certificate Details

Editing a Configuration

Publishing a Certificate

Withdrawing a Certificate

Deleting a Certificate

Note:
The actions allowable depend on the state of a certificate.

- If the certificate is not published, you can edit it, delete it, or publish it.
- If the certificate is published, you can withdraw it or edit it.
- If the certificate is withdrawn, you can edit it or delete it.

Creating a New Configuration

Note:
Creating a new configuration includes uploading a certificate to Control; however, this process is insecure because the certificate could be hijacked through a malicious browser extension. Limelight advises that you open the page in an incognito window to create the certificate because incognito pages block all browser extensions.

1. Click the **+ new** button at the top of the screen.
2. Fill out fields. All fields except *Intermediate certificates* are required. See Certificate Field Reference for details.
3. Click the **Create** button at the bottom of the screen.
4. The system verifies the contents of all uploaded files and displays errors if verification is unsuccessful.
5. If all fields pass validation:
   - the certificate is added to the list
   - the certificate's status ('New'), and version (v1) is added to the controls on the right side of the certificate's row along with the creation date and your user
6. Publish the certificate. See Publishing a Certificate.

Viewing Certificate Details

You can view additional certificate details that are not displayed in the list of certificates.

1. On the certificate list page, click the **View** icon for the certificate you want to examine.

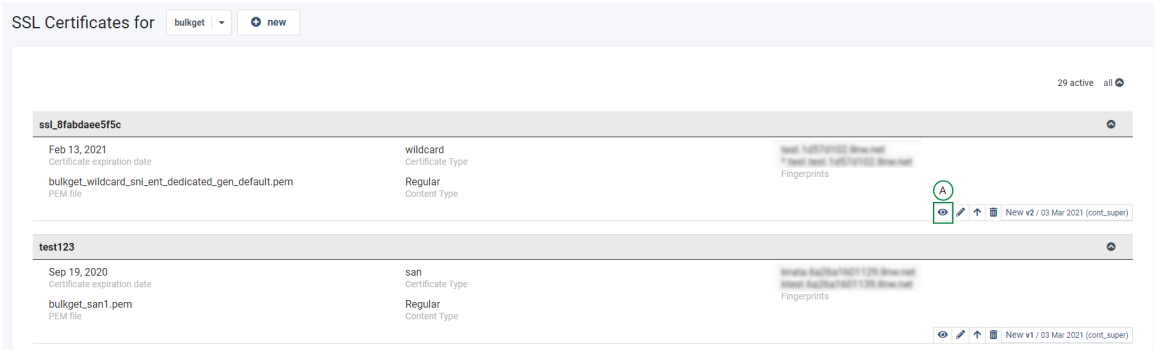

   In the screenshot, (A) is the **View** icon.

2. Details are displayed on a new page.

> **Note:**
> From this page, you can also take other options depending on the certificate status:
> - If the certificate is not published, you can edit it, delete it, or publish it.
> - If the certificate is published, you can withdraw it or edit it.
>
> See the following sections for instructions:
> - Editing a Certificate
> - Publishing a Certificate
> - Deleting a Certificate

Editing a Configuration

Use the edit icon to edit a configuration. Depending on the published state of the configuration, the edit icon has one of two tooltips:

- **Update/Rollback**: Modify a published configuration.
- **Update**: Modify a non-published configuration.

The two uses are essentially the same; both allow you to edit the configuration.

To edit a configuration:

1. Click the edit icon for the certificate you want to modify.

> **Note:** Although the Customer certificate and Certificate private key file fields are empty, the configuration defaults to the files contained in the configuration before you opened it in edit mode.

2. Make the desired modifications. See Certificate Field Reference for details.
3. Click the **Update** button.
4. If you changed any upload files, Control validates the files.
5. If all fields pass validation:
   - The changes are saved.
   - The certificate is added to the list.
   - The certificate's date and version are incremented.
   - A popup message is displayed, reminding you that although you have updated the configuration, you still need to publish it.
6. Publish the certificate. See Publishing a Certificate.

> **Note:** Each time you modify a configuration, its version is incremented.

Publishing a Certificate

When you publish a certificate, it gets pushed to the edge.

> **Note:** Control does not allow you to publish a certificate with domains covered by other published certificates. If you attempt to do so, you receive an error.

1. Click the **Publish** icon for the certificate you want to publish.

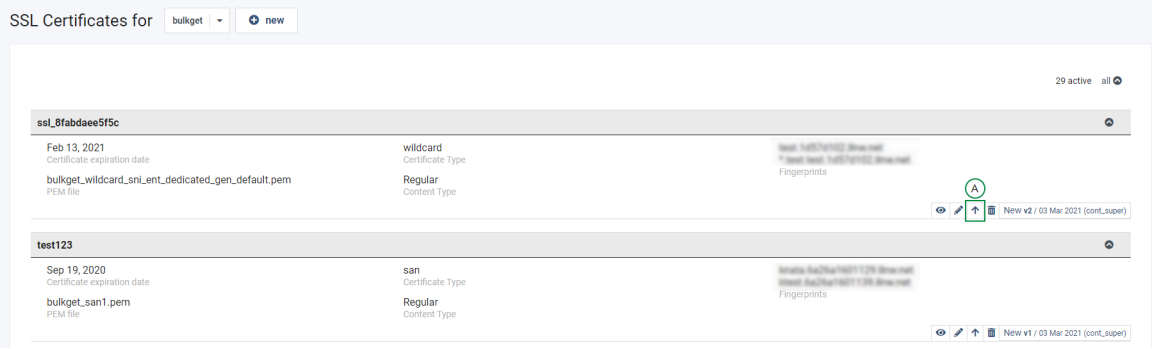

In the screenshot, (A) is the **Publish** icon.

2. Click **OK** in the dialog that prompts you to confirm.
3. The system starts a workflow that publishes the certificate to the edge. It normally takes about 6 hours to propagate changes.

Withdrawing a Certificate

When you withdraw a configuration it gets removed from the edge.

1. Click the **Withdraw** icon for the certificate you want to delete.

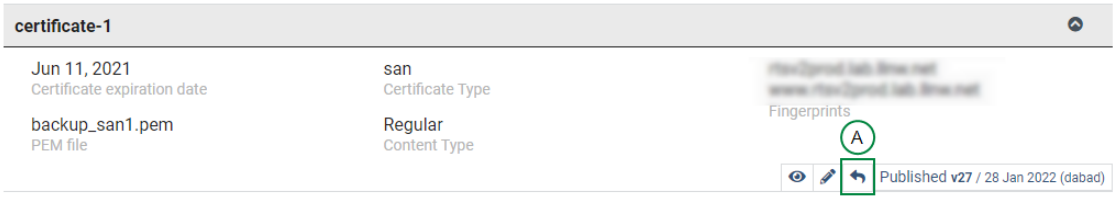

In the screenshot, (A) is the **Withdraw** icon.

2. Click **OK** in the dialog that prompts you to confirm.
3. The system starts a workflow that withdraws the certificate. It normally takes about 6 hours to propagate changes.

Deleting a Certificate

If a certificate has been published, you must withdraw it before you delete it. See Withdrawing a Certificate.

1. Click the **Delete** icon for the certificate you want to delete.

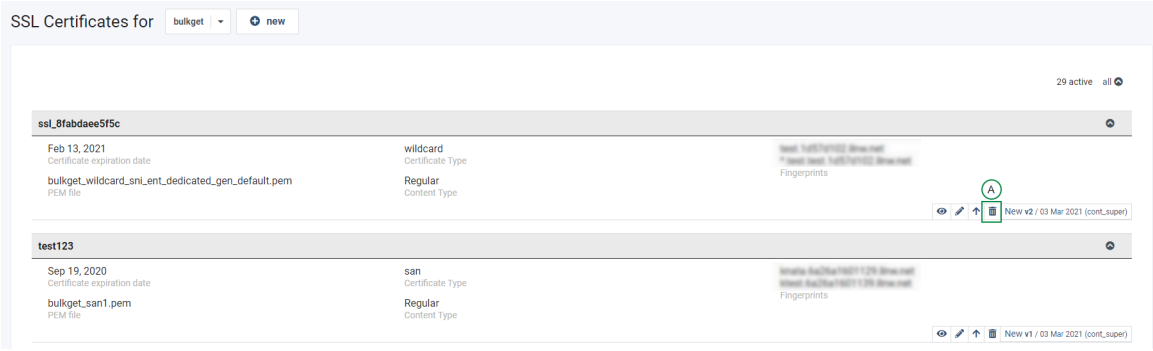

In the screenshot, (A) is the **Delete** icon.

2. In the resulting dialog, enter the certificate name as instructed, then click the **Delete** button.
3. The system deletes the certificate and starts a workflow that removes the certificate from the edge.

Certificate Field Reference

Field	Description
Certificate name	Name of the certificate to create; maximum 100 characters
Content Type	Type of content that the certificate encrypts; possible values: • **Regular** - Will use General Pool VIPs. Always select this option. • **LargeObject** • **Blue** **Note:** The value you select must match the collection of VIPs that can serve the shortname. For example, General Pool shortnames can only be served from General Pool VIPs. If you set a certificate on a shortname but you don't select 'Regular', the certificate will not work.

Field	Description
Customer certificate	Use this field to upload the certificate. Must be an X.589v3 ASCII Base64 PEM-encoded type and not password-protected. Customers that don't use that encoding must use Open SSL commands to convert their type to the required type. **Note:** This file almost always contains intermediate certificates. If not, use the **Intermediate certificate** field to upload the desired files.
Intermediate certificate	In most cases, this field should be left empty; intermediate certificate(s) will be derived from the provided certificate by default. Only upload the intermediate certificate(s) here if you know that the certificate encoding in the **Customer certificate** field does NOT already include intermediate certificate information.
Certificate private key file	Use this field to upload the private key. The key must be X.589v3 ASCII Base64 PEM-encoded and not password-protected.

Configuring Log Delivery Service

When requests for your content enter the CDN, the requests are logged based on Log Delivery Service configurations. The Log Delivery Service allows you to configure and manage your log files.

> **Note:**
> Log field names, delimiters, date and time format, file name, and directory structure adhere to W3C/ISO standards.

Log Delivery List Page

Navigate to Configure > Log Delivery Service in the navigation pane. The Log Delivery Service page is displayed and initially shows configurations for the account in the drop-down menu on the right above the list.

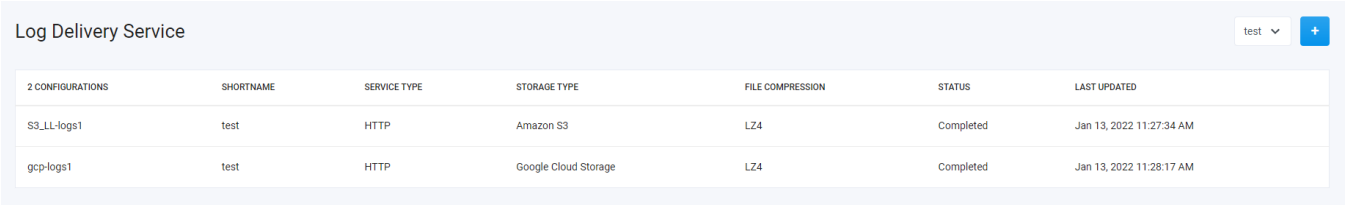

Each configuration in the list includes the following information:

Field	Description/Instructions
Number of configurations and configuration names	Customer-assigned configuration name.
SHORTNAME	Currently selected account name.
SERVICE TYPE	Delivery service for which logs will be created. HTTP is the only service supported.
STORAGE TYPE	Log file location. Possible values: • Origin Storage: Logs are stored at the root of your space in a directory called '_livelogs'. You are responsible for data maintenance; Limelight is not responsible for data removal. • Amazon S3: Amazon's cloud-based object storage. • Google Cloud Storage: Google's cloud-based object storage
FILE COMPRESSION	File compression method. Possible values: • ZSTD

Field	Description/Instructions
	• LZ4 • SNAPPY • LZF • GZIP **Note:** Limelight's LZ4 implementation uses the LZ4 (Framed) compression algorithm.
STATUS	Configuration status. When you create and save a configuration, it goes through a validation process. Possible status values: • In Progress • Pending • Completed • Failed • Deactivated
LAST UPDATED	Configuration's creation or last modified date.

Choosing an Account

Each account has its own set of Log Delivery Service configurations. You can choose an account to work with in the drop-down menu "A" on the right above the list.

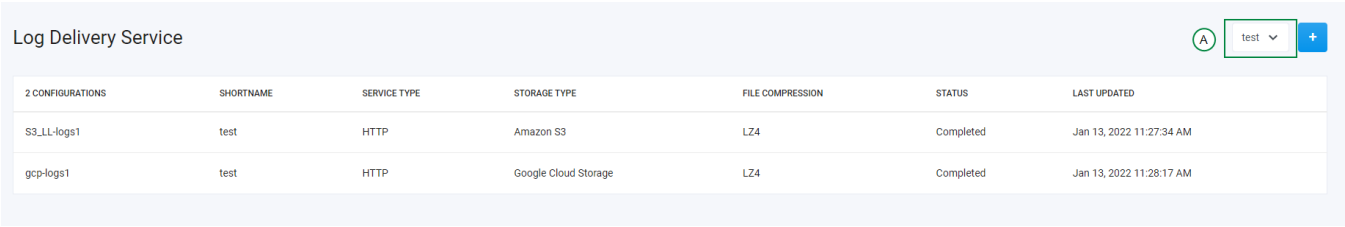

This list is more focused than the company/account drop-down menu at the top of the page and is limited to the accounts that your user can access and accounts that have the Log Delivery Service product enabled.

Working with Log Delivery Service Configurations

Creating a Log Delivery Configuration

Editing a Log Delivery Configuration

Configuring Log Fields

Deleting a Log Delivery Configuration

Deactivating a Log Delivery Configuration

Activating a Log Delivery Configuration

Enabling Log Delivery to Amazon S3

Enabling Log Delivery to Google Cloud Storage

[Enabling Log Delivery to Origin Storage]

Creating a Log Delivery Configuration

You can create a single configuration for any combination of shortname, storage location, and service type.

1. Click the **+** button at the top of the Log Delivery List Page.
   The *Add Configuration* page is displayed, and the following message warns you about extra fees if you choose to store logs in Limelight Origin Storage:

   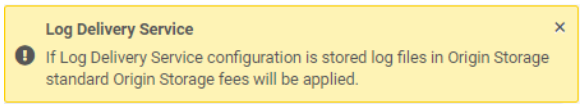

2. Fill out the fields at the top of the page, noting that required fields are marked with an asterisk in the user interface. See Log Delivery Configuration Fields for details.
3. Select fields to include in log files. See Configuring Log Fields.
4. Save the configuration by clicking the **Save** button.

> **Note:** It can take 15 to 40 minutes for a new configuration to take effect.

Editing a Log Delivery Configuration

1. Click the configuration's row on the Log Delivery List page.
   The configuration is displayed in edit mode. If the configuration's storage location is LimelightOrigin Storage, the following message is displayed that warns you about extra fees:

   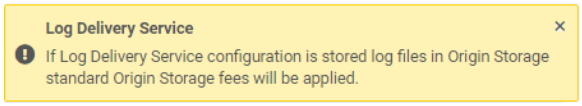

   > **Notes:**
   > - Existing configurations include Directory Layout and FileName Template Fields
   > - If your user does not have 'Manage' permissions for Log Delivery Service all fields are disabled and you cannot modify the configuration.

   .
2. Modify fields as needed. See Log Delivery Configuration Fields and Configuring Log Fields for details.
3. Save the configuration by clicking the **Save** button.

> **Notes:**
> - It can take 15 to 40 minutes for changes to take effect.
> - Depending on your Log Delivery Service permissions, you may not be able to edit a configuration.

Configuring Log Fields

You can add, remove, and reorder active log fields. You can also add static fields.

Moving Fields between Lists

- Drag and drop individual fields from one set to another.
- Move all fields using the button beneath the **Selected log fields** set.
  - Click **SELECT ALL** to move all fields from the **Available log fields** set to the **Selected log fields** set. The button's text changes to 'DESELECT ALL'.
  - Click **DESELECT ALL** to move all fields from the **Selected** set to the **Available** set. The button's text changes to 'SELECT ALL'.

Reordering Selected Fields

Drag and drop individual fields to reorder them.

Working with Static Fields

Static fields are user-defined fields with a value that does not change.

- To add a static field:
  1. Click the **ADD STATIC FIELD** button; then enter a field name and value in the subsequent dialog,
  2. Click **ADD ACTIVE FIELD**.

  The field is added to the **Available log fields** set. From there you can move it to the **Selected log fields** set.
- To edit or delete a static field:
  1. Click the field.
  2. In the subsequent dialog enter a new value and click **SAVE**, or click the **DELETE** button.

Deleting a Log Delivery Configuration

1. Click the configuration's row in the Log Delivery List page.
   The configuration is displayed.
2. Click the **DELETE** button at the bottom of the page.
3. Agree to the deletion in the subsequent confirmation dialog.
   Control deletes the configuration.

> **Note:** It can take 15 to 40 minutes for the deletion to take effect.

Deactivating a Log Delivery Service Configuration

You can deactivate a Log Delivery Service configuration for purposes such as forcing the configuration to stop gathering log data.

1. Click the configuration's row in the Log Delivery List page.
   The configuration is displayed.
2. Click the **DEACTIVATE** button at the bottom of the page.
   A confirmation message is displayed at the top right of the page and the button's label changes to **ACTIVATE**..
   The configuration's status on the Log Delivery List page changes to **Deactivated**.

> **Note:** It can take 5 to 10 minutes for a deactivation to take effect.

Activating a Log Delivery Service Configuration

You can reactivate a deactivated Log Delivery Service configuration.

1. Click the configuration's row in the Log Delivery List page.
   The configuration is displayed.
2. Click the **ACTIVATE** button at the bottom of the page.
   A confirmation message is displayed at the top right of the page and the button's label changes to **DEACTIVATE**..
   The configuration's status on the Log Delivery List page changes to the state it was in before it was deactivated.

> **Note:** It can take 15 to 40 minutes for an activation to take effect.

Enabling Log Delivery to Amazon S3

You can store your log files on the Amazon S3 platform. Amazon S3 is a cloud object storage service built to store and retrieve data.

Prerequisites

Before configuring Amazon S3 as a storage location, you must do the following:

- Create an S3 Identity and Access Management (IAM) user in Amazon's configuration screens.
- Give the IAM user the following permissions for the bucket where you want to store logs:
  - ListBucket
  - GetObject
  - PutObject

Configuration Fields

These are visible only when you select Amazon S3 as the storage location.

Field	Description
REGION	S3 bucket geographic area.
BUCKET NAME	S3 bucket title.
PATH	Path within bucket where logs are stored. **Note:** Do not add a leading slash to the path. If you do, Amazon creates an object URL with a double slash. Example: `https://bucket.s3.region.amazonaws.com//cdn_logs...`
ACCESS KEY	Bucket access key provided by Amazon.
SECRET KEY	Bucket secret key provided by Amazon. **Note:** After you set the secret key and save the configuration, the key is not visible, but you can enter a new key if needed and save the configuration.

Enabling Log Delivery to Google Cloud Storage

You can store your log files on the Google Cloud Storage platform. Google Cloud Storage is a service for storing and accessing your data on Google Cloud Platform infrastructure.

Prerequisites

Before configuring Google Cloud Storage as a storage location, you must do the following:

1. Create a Google Cloud Project (GCP) or use an existing project. See Google's Google Cloud Platform - Creating and managing projects guide for instructions.
2. Set up a GCP bucket to store your logs. You can create a new bucket or use an existing one. See Google's Create storage buckets guide for instructions..
3. Create a Google service account that Log Delivery Service will use to access your bucket. See Google's Service accounts guide for instructions.
4. Using Google's IAM roles for Cloud Storage, guide, grant the following roles on the bucket:
   - Storage Object Creator (`storage.objectCreate`)
   - Storage Object Viewer (`storage.objectViewer`)
5. Add the service account as a member of the bucket you created in step 2.
6. Generate JSON access keys for the service account. See Google's Creating service account keys guide for instructions.

Configuring a Google Cloud Storage Location

1. Select **Google Cloud Storage** in the **STORAGE LOCATION**drop-down menu.
2. Configure the fields described in Configuration Fields.
3. Click **SAVE**.

Configuration Fields

These are visible only when you select Google Cloud Storage as the storage location. Required fields are marked with an asterisk in the Control user interface.

Field	Description
CLIENT EMAIL	Value of the `client_email` field in the JSON file associated with the Google service account you created.
SECRET KEY	Value of the `private_key` field in the JSON file associated with the Google service account you created. **Note:** After you set the secret key and save the configuration, the key is not visible, but you can enter a new key if needed and save the configuration.
BUCKET NAME	Title of the Google Cloud Storage bucket you created.
PATH	Path within the bucket where logs are stored. Defaults to an empty value.

Field	Description
	Note: Do not add a leading slash to the path. If you do, Google Cloud Storage creates an object URL with a double slash. Example: `gs://bucket_name//cdn_logs/...`

Enabling Log Delivery to Origin Storage

You can store your log files on the Origin Storage platform. Origin Storage is a distributed storage service operated by Limelight Networks.

Note: Standard fees apply for using Origin Storage.

Prerequisites

Origin Storage must be enabled for the name selected in the **SHORTNAME**drop-down menu.

Configuring the Origin Storage Location

1. Select **Origin Storage** in the **STORAGE LOCATION** drop-down menu.
2. Configure the field described in Origin Storage Configuration Field.
3. Click **SAVE**.
   If Origin Storage is not enabled for the selected shortname, the following message is displayed when you attempt to save the configuration:

   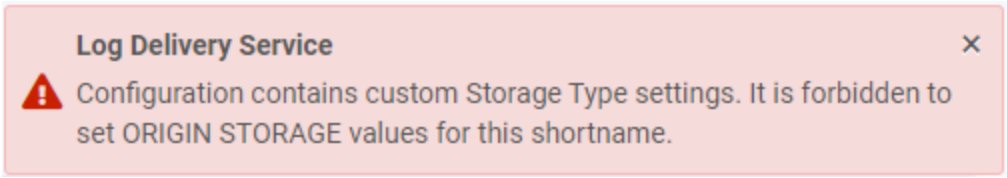

   Contact your Limelight Account Manager to enable Origin Storage for the shortname.

Origin Storage Configuration Field

Field	Description
STORAGE ACCOUNTS	The Origin Storage account where you want to store logs. By default logs are stored under the same account that owns LDS configuration

Working with Personally Identifiable Information

Limelight's Log Delivery Service conforms to General Data Protection Regulations (GDPR) requirements.

You can configure logs to include the following fields, which contain Personally Identifiable Information (PII) :

- cs-cookie
- cs-uri
- so-src-uri

Signing PII Agreements

Per GDPR, you must explicitly indicate that you understand risks associated with the PII fields.

When you access Log Delivery Service, a message that describes the risks involved is displayed:

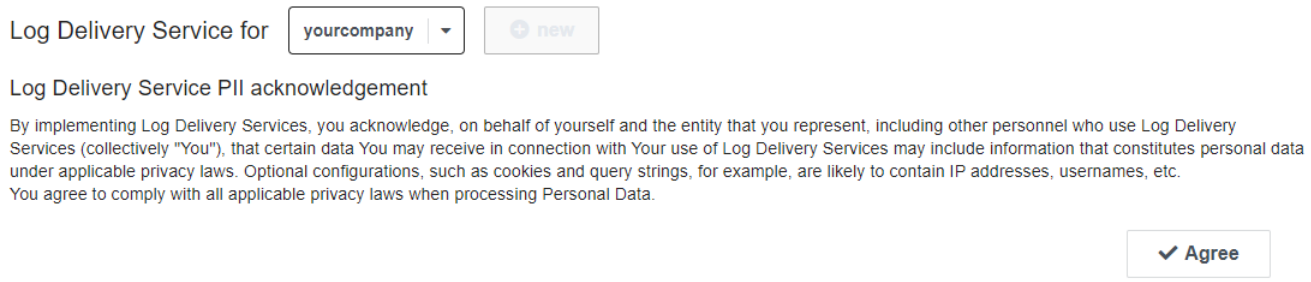

Click the **Agree** button to indicate you agree.

Notes:

- If you do not agree to the terms and conditions, you cannot view any Log Delivery Service configurations.
- Non-Company Admin users can sign agreements only for the company to which they belong.
- Company Admin users can sign agreements for child companies as well.

Field Reference

Log Delivery Service Configuration Fields

Delivery Destination Fields

Delivery Options Fields

Log File Fields

Log Delivery Service Configuration Fields

Note: Log Delivery Service configuration fields are attributes of a Log Delivery Service configuration and are not to be confused with log fields (see Log File Fields), which appear in log files.

Field or Section	Description
CONFIGURATION NAME	Customer-assigned configuration name.
SHORTNAME	The shortname to which the configuration applies.
SERVICE TYPE	Delivery service for which logs will be produced. Only HTTP is available for selection.
Delivery Destination	See Delivery Destination Fields.
Delivery Options	See Delivery Options Fields.

Delivery Destination Fields

Field or Section	Description
STORAGE TYPE	Log file location. Possible values: • Origin Storage: Logs are stored at the root of your space in a directory called '_livelogs'. You are responsible for data maintenance; Limelight is not responsible for data removal. • Amazon S3: Amazon's Simple Storage Service. • Google Cloud Storage: Google's cloud-based object storage. See Google Cloud Platform - Creating and managing projects. **Note:** If you change the location from Amazon to Origin Storage, the following message is displayed: 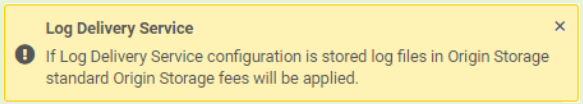
STORAGE ACCOUNT	The Origin Storage account where you want to store logs. By default logs are stored under the same account that owns the Log Delivery Service configuration.

Delivery Options Fields

Field	Description

DIRECTORY LAYOUT	The directory structure where your log files are stored. Read-only.
	`/http/{config_uuid}/YYYY/MM/DD`
	EdgeQuery creates the directory replacing `config_uuid` and `YYYY/MM/DD` with the correct values.
	
	Descriptions of Directory Components
	`http` - service name of the request.
	`{config_uuid}` - universally unique ID of a log delivery configuration.
	`YYYY/MM/DD` - the date when the request completed. All content requests completed on a given YYYY/MM/DD date are stored in a log in the DD directory.
	Example: The `config_uuid` is `config_1`, and the requests were completed on 2020/08/01 and 2020/08/02.
	`config_1`
	`|`
	`+--2020`

Field	Description

	<pre>| |
	| +--08
	| | |
	| | +--01
	| | | |
	| | | +--<log file stored here>
	| | |
	| | +--02
	| | | |
	| | | +--<another Log file</pre>

FILE NAME TEMPLATE	Log file naming convention. Read-only. `{shortname}_{request_end_date_time_from}-{request_end_date_time_to}.{process_window_date_time}_{split_id}.{format}.{compression}` Descriptions of Filename Components `{shortname}` - account name. `{request_end_date_time_from}`, `{request_end_date_time_to}` - the start and end time for requests contained in the log. `{process_window_date_time}` - the date/time when a batch of log files were processed (prepared) for a customer. `{split_id}` - ID of the file if a log file needs to be split to avoid large files. Each file that is part of a split is assigned an ID. The first file that is part of a split contains 000 for `split_id`; subsequent files contain `001`, `002`, and so on. **Notes:** `split_id` is always present in the file name. If a file does not need to be split, the value is 000. The log file size limit is 1GB. A new file with a new *split_id* will be created once the file reaches 1GB. Log file size is measured before compression, so a log file may be split even though the size is smaller than 1GB. All strings in log file names are URL-encoded. `{format}` - file naming specification. Always `w3c`. `{compression}` - compression type.
FILE COMPRESSION	File compression method. **Note:** Limelight's LZ4 implementation uses the LZ4 (Framed) compression algorithm.

Field	Description
	Note: Limelight encourages you to investigate available compression methods before deciding on a method.

Log File Fields

The following fields are available for you to include and order in your logs.

Field	Description / Instructions
c-asn	The autonomous system number calculated based on client IP address.
c-city	The City name derived from the client IP address using the IPGeo DB.
c-country	The Country name derived from the client IP address using the IPGeo DB.
c-ip	The Client IP Address (end-user).
c-port	The client remote port number used for a connection.
c-state	The State name derived from the client IP address using the IPGeo DB.
cs-accept-language	The value of the Accept-Language request header.
cs-cmcd	The CMCD metric sent by a compatible chunk streaming media player as specified by CTA-5004 saved in query term URL-encoded format, regardless of methods used to ingest by player.
cs-cookie	The URL-encoded cookie HTTP request header. GDPR Personally Identifiable information is included.
cs-custom-header1, cs-custom-header2, cs-custom-header3, cs-custom-header4, cs-custom-header5	The value of the request header specified in the `log_request_header` rewrite option. You can include the value of up to five custom headers as defined as `log_request_header*` fields in Caching & Delivery (v2). The log_request_header* fields correspond to cs-custom-header* as follows: log_request_header - cs-custom-header1 log_request_header2 - cs-custom-header2

Field	Description / Instructions
	log_request_header3 - cs-custom-header3 log_request_header4 - cs-custom-header4 log_request_header5 - cs-custom-header5 For example, to include cs-custom-header1, do the following: **Step 1**. Edit a configuration in 'Configure > Caching & Delivery (v2)'. Go to 'Add custom request header' within the 'Headers & Methods' section. Enter: • 'log_request_header' in the 'Header name' field • a value for 'log_request_header' in the 'Header value' field. **Step 2**. In 'Configure > Log Delivery Service', add 'cs-custom-header1' to the 'Active List'. Then save the configuration.
cs-headers	The value of the HTTP request headers specified in the `log_req_header` rewrite option. These headers are logged as key-value pairs in this field. If multiple headers are specified to be logged, each key-value pair is separated by a comma. The maximum size of this field is 2048 bytes. If the maximum size is exceeded, `error=toolarge` is logged.
cs-http-proto	The version of the HTTP protocol sent from the client to the server.
cs-method	The HTTP request method (GET, POST, and so on) sent from the client to the server.
cs-range	The value of the Range header sent from the client to the server. URL-encoded.
cs-referer	The value of the Referrer header sent from the client to the server. URL-encoded.
cs-ssl-cipher	The version that the client supports, sent from the client to the server.
cs-ssl-proto	The version that the client supports, sent from the client to the server.
cs-uri-host	The domain part of the Published URL.
cs-uri-noquery	The URL-encoded published URL (query part excluded).

Field	Description / Instructions
cs-uri	The URL-encoded published URL that includes query strings. Includes GDPR Personally identifiable information.
cs-user-agent	The value of the User-Agent header in the request from the client to the server. URL-encoded.
date	The request end time (date part) in *yyyy-MM-dd format* (UTC time zone).
datetime	The request end time in yyyyMMddHHmmss format (UTC time zone).
duration	The request duration in milliseconds.
s-dest-addr	The IP address that the end user connects to. It is most often a virtual IP associated with a request router. In rare cases, when alternative request routing is configured, this IP address corresponds directly to a caching server.
s-host	The hostname of the Limelight server that received the request.
s-ip	The IP address of the edge-most Limelight server that received the request.
s-pop	The Limelight PoP name of the server that received the request.
s-ttfb	The number of milliseconds between the CDN receiving the end-user request and writing the first byte of the response, as measured on the server. A value of 0 (zero) means the time was less than 1ms.
sc-bytes	The number of response bytes, modified to include the packet and retransmit overhead.
sc-content-length	The value of the Content-Length header in the response from the server to the client.
sc-content-type	The value of the Content-Type header in the response from the server to the client.
sc-headers	The value of HTTP response headers specified in the `log_resp_header` rewrite option. These headers are logged as key-value pairs in this field. If multiple headers are specified to be

Field	Description / Instructions
	logged, each key-value pair is separated by a comma. The maximum size of this field is 2048 bytes. If the maximum size is exceeded, `error=toolarge` is logged.
sc-request-id	The unique ID that identifies a request (generated by the server and sent to the client in the `X-LLNW-Dbg-Request-Id` response debug header)
sc-rexb	The number of bytes retransmitted in the response from the server to the client.
sc-rtt	The client socket smoothed round-trip time in microseconds.
sc-rttv	The client socket smoothed round-trip time variance in microseconds.
sc-status	The HTTP status code in the response from the server to the client. In addition to standard Content Delivery status codes, the sc-status field may contain non-standard status codes: • 000 - A Limelight-specific status code returned when the origin sends no response, so there is no status code to log (for example when the client disconnects before the origin delivers the response). • 600 - A Limelight-specific status code indicating the origin returned a non-HTTP-compliant response so a status code could not be obtained. For a list of standard status codes, see the 'Response Codes' section in the Content Delivery User Guide.
so-src-uri-noquery	The URL-encoded source/ origin URL that the published URL has been mapped to (query part excluded).
so-src-uri	The URL-encoded source/ origin URL that the published URL has been mapped to.
time	The request end time (time part) in *HH:mm:ss.SSS* format (UTC time zone).
x-firstnode-cached	Integer value indicating whether a cache hit occurred on the Limelight server that received the request Possible values: 0 - a cache miss occurred

Field	Description / Instructions
	1 - a cache hit occurred Customers can use the field to calculate cache efficiency in terms of requests. **Note:** This field reflects a hit or miss on only the first cache node involved. It does not reflect cache hits and misses for the entire CDN.
x-log-key-value	The string representation of the key value pairs configured via the log_keyval rewrite option, the ArcLight `llnw.log_keyval() builtin`, and the `log_keyval_header` global option. This column is limited to 1024 bytes. Limelight configures the EdgePrism key-value pairs on behalf of customers. Please contact your Account Manager if you are interested in this feature.

Retrieving Log Files from Origin Storage

You can retrieve your files from Limelight Origin Storage using Origin Storage API calls in conjunction with an HTTP GET request or via the Origin Storage Management Console.

Download Using the API

Download Using the Storage Management Console

Download Using the API

All methods in this section are in the Origin StorageJSON-RPC API interface. We presented essential information; for detailed information about each method, see the Origin Storage API Reference Guide.

Introduction to Methods

This section describes the methods you need to download files.

Obtain an Origin Storage Token

List Log Files

Obtain a Protected Download URL

API End-to-End Example

Obtain an Origin Storage Token

Use the `login` method available in the Origin Storage JSON-RPC interface. The token string that allows you to make authenticated calls in the JSON-RPC interface. There are several methods of logging in, but we will use the simplest.

login Signature

`login( username, password, detail)`

Parameters

`username`: Your API user name.

`password`: Your API user name.

`detail`: A boolean indicating whether you want simple data or more extensive data returned.

List Log Files

To list log files, call the `listFile` method available in the Origin Storage JSON-RPC interface.

listFile Signature

`listFile( token, dir, pageSize, cookie, stat)`

Parameters

`token`: The token returned from the `login` call.

`dir`: A string representing the directory for which you want a list of files.

`pageSize`: A number indicating the number of results (files) to return.

`cookie`: A number used for making multiple `listFile` calls for paginated results.

`stat`: A boolean whether to include file details.

Obtain a Protected Download URL

To eliminate security risks, you can obtain a time-based URL to download your log files. This is the `mediaVaultUrl` method available in the Origin Storage JSON-RPC interface.

First, use the `mediaVaultUrl` method to obtain a secure download URL, and then use an HTTP GET request to download.

mediaVaultUrl Signature

`mediaVaultUrl(token, path, expiry)`

Parameters

`token`: The token returned from the `login` call.

`path`: File to generate MediaVault URL.

`expiry`: Download URL expiry for an object in seconds.

The method returns this object:

```
{
  "code": 0,
  "download_url": "http://cs-download.limelight.com/<path to file>",
  "message": "success",
  "preview_url": "http://cs-download.limelight.com/<path to file>",
}
```

Note:
Do not attempt to directly download content from Origin Storage using FTP, SFTP, FTPS, SCP, or rsync because doing so can negatively impact other system processes. To download content, use an HTTP GET request.

API End-to-End Example

For simplicity, we've omitted error checking. The code sample uses Python.

```
import jsonrpclib
import requests

url = 'http://{Account name}.upload.llnw.net/jsonrpc'
api = jsonrpclib.Server( url )
res = api.login(yourUser, yourPassword, True)
token = res[0]

'''
User-defined variables
'''
storage_log_dir = '/{account name}/_livelogs/'
pageSize = 10000  # page size for listing log files
files_to_download = []  # log files to download
media_vault_expiry = 60  # expiry time for mediaVaultUrl request
mv_errors = {-1: "Internal error", -2: "Path exists and is a directory", -8:
"Invalid path",
            -34: "Invalid expiry", -60: "Service is disabled or unavailable", -
10001: "Invalid token"}
'''
Function to examine files returned from calls to listFile
Based on a condition that you determine, you write file names to a list
of files that will later be downloaded.
This simple example looks for file names that contain the number 2.
'''
def parse_list(file_list):
    for (log_file) in file_list:
    name = log_file['name']
    if name.find('2') > -1:
        files_to_download.append(name)
        print(log_file['name'])

'''
List Log files. This is a simplistic approach for demonstration purposes.
Customers might want to try a multi-threaded approach because the number of files
can be quite large
```

```
'''
results = api.listFile(token, storage_log_dir, pageSize, 0, True)
file_list = results['list']
if len(file_list) > 0:
    parse_list(file_list)
    cookie = results['cookie']
    while cookie > 0:
        results = api.listFile(token, storage_log_dir, pageSize, cookie, True)
        file_list = results['list']
        parse_list(file_list)
        cookie = results['cookie']

'''
Download file. This is a simplistic approach for demonstration purposes.
Customers might want to try a multi-threaded approach for a large number of files to
download.
'''
for file_name in files_to_download:
    log_path = storage_log_dir + '/' + file_name
    mvu = api.mediaVaultUrl(token, log_path, media_vault_expiry)
    if mvu['code'] != 0:
        print("Error attempting to call 'mediaVaultUrl.\nCode: " + str(mvu['code'])
+ ": " + mv_errors[mvu['code']])
        mv_download_url = mvu['download_url']
        # Use the requests library to make the download
        response = requests.get(mv_download_url)
        # Upon non-success write a line to your errors file
        if response.status_code != 200:
            print("Unable to download " + file_name + ". Status code: " +
response.status_code)
```

Download Using the Storage Management Console

You can download a log file using the Origin Storage Management Console.

Begin by logging into the Limelight Control Portal, then follow these steps:

1. Select "Manage", followed by "Origin Storage Console."
2. Navigate to the folder that contains the file you want to download.
3. Click the **download icon**.

Your browser downloads the file.

Configuring Live Streaming

Overview

You can access the Limelight Control portal to view and configure the slots you have purchased. You can view, create, copy, and more.

Main Configuration Page

Navigate to Configure > Live Streaming in the navigation pane.

The live slots configured for your specific shortname are listed. At the top of the page is summary information for the type of slots you have purchased. For the transcode and transmux slots, the display shows how many you have configured (used) and how many of each type you still have left to configure (available). All slot types are displayed regardless of whether you have purchased a particular slot type.

- If 0 slots are available, you cannot create any more of that specific slot type.
- If 0 slots are available and 0 are used, you have not purchased that type of slot.

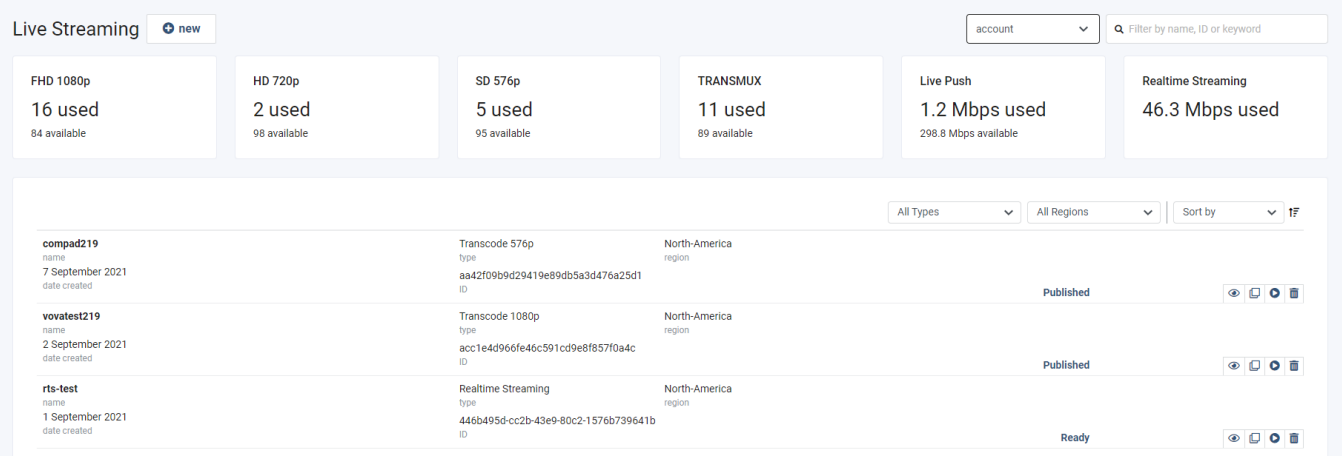

The display for Realtime Streaming and Live Push slots shows slightly different information:

- Realtime Streaming slots - the total number of slots used and the total of all ingest bitrates provisioned for all Realtime Streaming slots.
- Live Push - the total of all ingest bitrates provisioned, the total ingest bitrates available across all Live Push slots.

For more information about Realtime Streaming slots, please see the Realtime Streaming Guide.

Buttons and Icons

+ new button: At the top right of the screen, allows you to configure one of your purchased live slots.

Icons at the right side of each slot row allow you to take actions on the slot.

> **Note:** Icons available depend on the slot type.

Icon	Description
Edit	Lets you modify the configuration. See Edit a Slot.

Icon	Description
Details	Lets you view a specific live slot in detail. See View Slot Details. (You can also view slot details by clicking in the slot's row.)
Clone	Makes a copy of an existing live slot configuration. See Clone a Slot.
Record and Schedule	Specific to Live to VoD functionality.
Player	Allows you to view the live stream. See Viewing a Slot's Live Stream.
Delete	Allows you to delete a live slot configuration. See Delete a Slot.

List Information

Each entry in the list has information in the following table:

Item	Description
name	The name given to the slot when you created it.
date created	The date the slot was created.
type	The type of slot.
ID	The specific ID number given to your slot for tracking and routing purposes.
region	The geographic region into which the slot ingests.
status	The slot state. Possible values are in the following table:

State	Description
LiveEventProvisioning	The slot configuration is in the process of being set up. It can take several minutes for all the configurations to be enabled.
Not_Provisioned, Error	The configuration has not been set up on the ingest. Please contact Customer Support.
Published, Ready	The slot has been properly configured and is ready to use.

The status of the slot is shown to the left of the icons. Possible statuses are described in the following table:

Status	Description
LiveEventProvisioning	The slot configuration is in the process of being set up. It can take several minutes for all the configurations to be enabled.
Not_Provisioned,	The configuration has not been set up on the ingest. Please contact

Status	Description
Error	Customer Support.
Published, Ready	The slot has been properly configured and is ready to use.
Deleting	The slot is in the process of being removed.

Click on a specific slot (or click the **Details** icon) to view slot details.

Filtering and Sorting the List of Slots

Filtering, sorting, and sort direction controls appear above the list of slots:

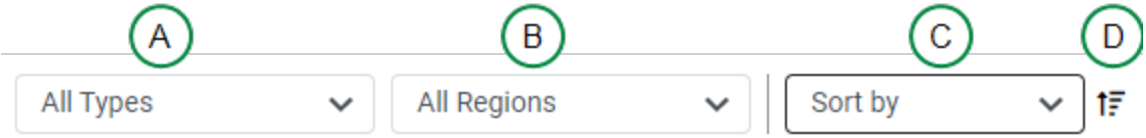

A - Filter by slot type

B - Filter by region

C - Sort by various fields

D - Sort direction

Filtering

Make selections in the Filter by slot type and Filter by region filters.

> **Note:** A region is the geographic area into which the slot ingests.

Sorting

By default, the list is sorted by date created. The arrow in the sort direction field points up, indicating the list is sorted in descending order.

Choose a field to sort by and a sort direction.

Configuring a Slot

Configuring a slot is a simple process.

1. Begin by clicking the **new** button at the top of the page and selecting a slot type:

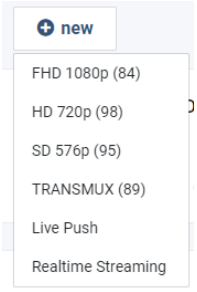

> **Note:** According to your account, the number in parentheses beside the transcode and transmux slots indicates the number of remaining slots you can create of each type.

2. Then, fill out fields on the following sections in the page that appears:
   Identifying information
   Ingest details
   Configuration Details
   Encoding details
   Content security
3. Review your configuration, making any corrections.
4. When finished, click the **Submit** button.

Identifying information

> **Note:** For information about fields specific to Live Push slots, please see the Video Live Push Guide.

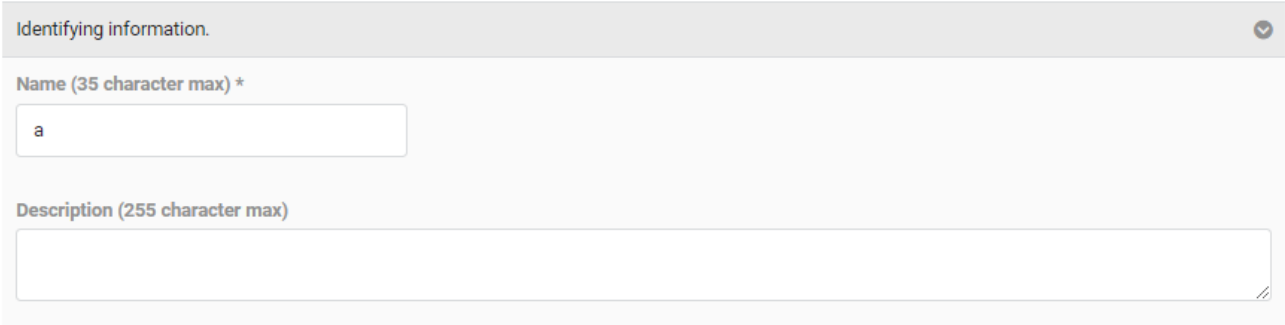

> **Note:** The fields available for configuration depend on the slot type. The following table describes fields for all slot types.

Field Name	Description	Notes / Instructions
Name	Identifies your slot.	A unique slot name is required; no two slots may have the same name. A name can only contain letters,

Field Name	Description	Notes / Instructions
		numbers, and hyphens (-). It cannot start or end with a hyphen. Spaces are not allowed.
Keywords, comma-separated	Optional keywords to tag your slot.	• To enter a tag, type the value then press the Enter key. The following characters are not allowed: ◦ period ('.') ◦ apostrophe (''') ◦ slash ('/') ◦ backward slash ('\') ◦ left bracket ('[') ◦ right bracket (']')
Description	An optional free-form field to describe your slot.	
Callback URL	The URL to which you want Limelight to publish events. See the 'Event Callback API Developers Reference' for more details.	

Ingest details

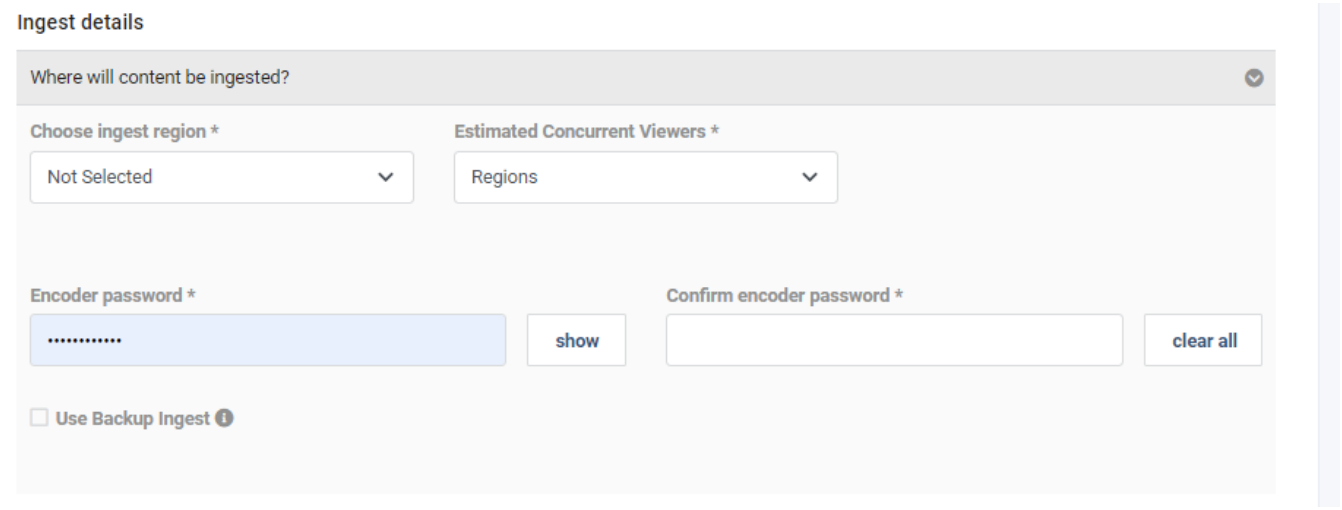

Note: The fields available for configuration depend on the slot type. The following table describes fields for all slot types.

Field Name	Description	Notes / Instructions
Choose ingest region	Region in which you want to ingest into your slot.	Select a region where your encoder is located or the region closest to your encoder's physical location.
Estimated Concurrent Viewers	Estimated concurrent viewers of the slot per region. Helps the Limelight CDN optimize output from your slot.	Select one or more regions. Each region you select causes a new field to appear in which you can enter the estimated number of viewers for the region.
Primary POP, Backup POP	Once you select a region, the available PoPs are listed in these drop-downs.	Select the primary that is closest to the physical location of your encoder to minimize data transmission time. Your backup selection must be different from your primary to provide redundancy for your live stream if there is a problem streaming to the primary. **Note:** The Backup POP field is only displayed if you check the 'Use Backup Ingest' checkbox.
Encoder password	Password that you configured for your stream while encoding it.	Enter and confirm the encoder password. Toggle password visibility with the **show**/**hide** button. Clear fields using the **clear all** button.
Use Backup Ingest	Stream to the backup ingest server as well as the primary.	Creating a backup allows you to stream to two distinctive ingest locations. When streaming to both primary and backup, you have redundancy that helps protect from regional outages and certain maintenance events. **Note:** When you select this option, the 'Backup POP' field is displayed.
Override Authentication Credentials	You can override the global password on a per-slot basis if desired. Configure your user and password by selecting this option.	Each slot uses the global user and password configured at the time your account was set up. See your Account Activation email for details.

Configuration Details

The Configuration Details section is visible only for Live Push slots.

Field Name	Description/Instructions	Notes
Total Ingest Bitrate	The desired bitrate for the slot.	The total ingest bitrate in kilobits per second.
Chunk Duration	The desired durations for HLS and DASH formats.	Larger values result in your player making fewer requests. Lower values result in lower latency, but the number of chunk requests increases and may result in higher traffic costs and potential playback issues.
File Expiry Time	Duration (in seconds) to keep the content on the Live Push Ingest server.	Larger values result in files being available longer but also use more storage space.
Publish IP Access	A list of IP addresses allowed to access the Live Push stream.	You can enter a single IP address or a range. Ranges must be input in CIDR format; for example, for a range of `10.10.10.0` - `10.10.10.255`, the input value should be `10.10.10.0/24`.
Username	Username for Live Push slot access.	Each Live Push slot requires at least one set of authentication credentials. Multiple sets of credentials can be configured by clicking **+ add** after entering each set of credentials.
Password	Password for Live Push slot user access. Click **Show** to display the password.	Passwords should be strong.
Confirm Password	Confirmation of the password.	Must match the password.
Publish Credentials	A list of the credentials to publish to enable access to the Live Push slot for listed users. To add a credential set: 1. Click the **add** button. The *ADD PUBLISH*	The listed users will have access to the Live Push slot.

Field Name	Description/Instructions	Notes
	CREDENTIALS dialog is displayed. 2. Enter a user name and a password, then confirm the password by entering it in the **Repeat Password** field. 3. Click the **Add** button. The credentials are added beneath the **add** and **clear all** buttons. To clear a single credential set, hover the mouse pointer over the row and click **remove**. To clear all credential sets, click the **clear all** button.	

Encoding details

Encoding detail options are largely similar for all slot types with a few differences.

Transcode Slots

Transmux Slots

Realtime Streaming Slots

Subtitles and Timecodes

Note: The Encoding Details section is not available for Live Push streams.

Note: For Realtime Streaming Encoding Details, please see 'Creating a New Slot Configuration' in the Realtime Streaming Guide.

Transcode Slots

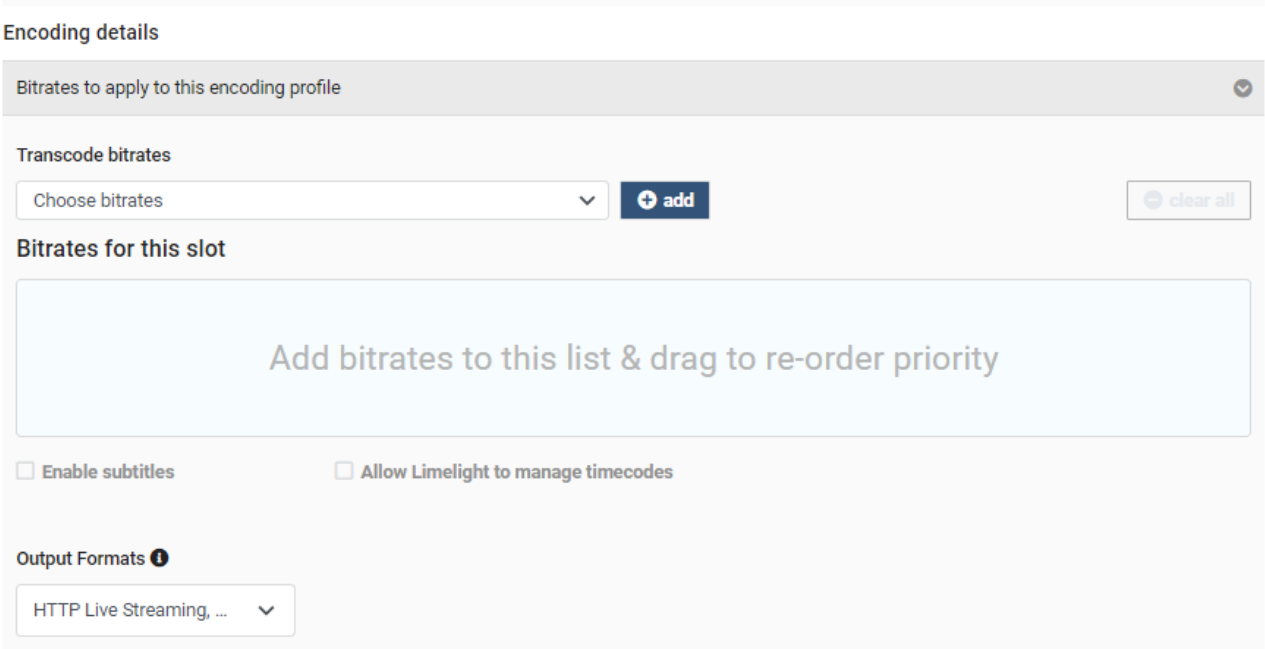

Field Name	Description/Instructions	Notes
Transcode bitrates	The bitrates you want to output from the slot. Choose one or more bitrates from the drop-down menu, then click the **add** button to add the bitrates to the 'Bitrates for this slot' field.	Each selection is an encoding profile All configured bitrates must be published by your encoder to MMD Live for the slot to function correctly.
Bitrates for this slot	Displays all selected bitrates. If you have 'Bitrate Order' enabled for your account, a drag handle displays when you hover your mouse to the left the slot's order number. Use the drag handle to drag and drop bitrates to reorder them. The order number determines the order in which the bitrate URL appears in the output manifest file.	MSS output will not support a custom order in which the audio-only bitrate is placed first. To remove a bitrate, hover over its row and click the **remove** button.
• **Enable subtitles** • **Allow Limelight to manage timecodes** • **Output Formats**	See Subtitles and Timecodes.	

Transmux Slots

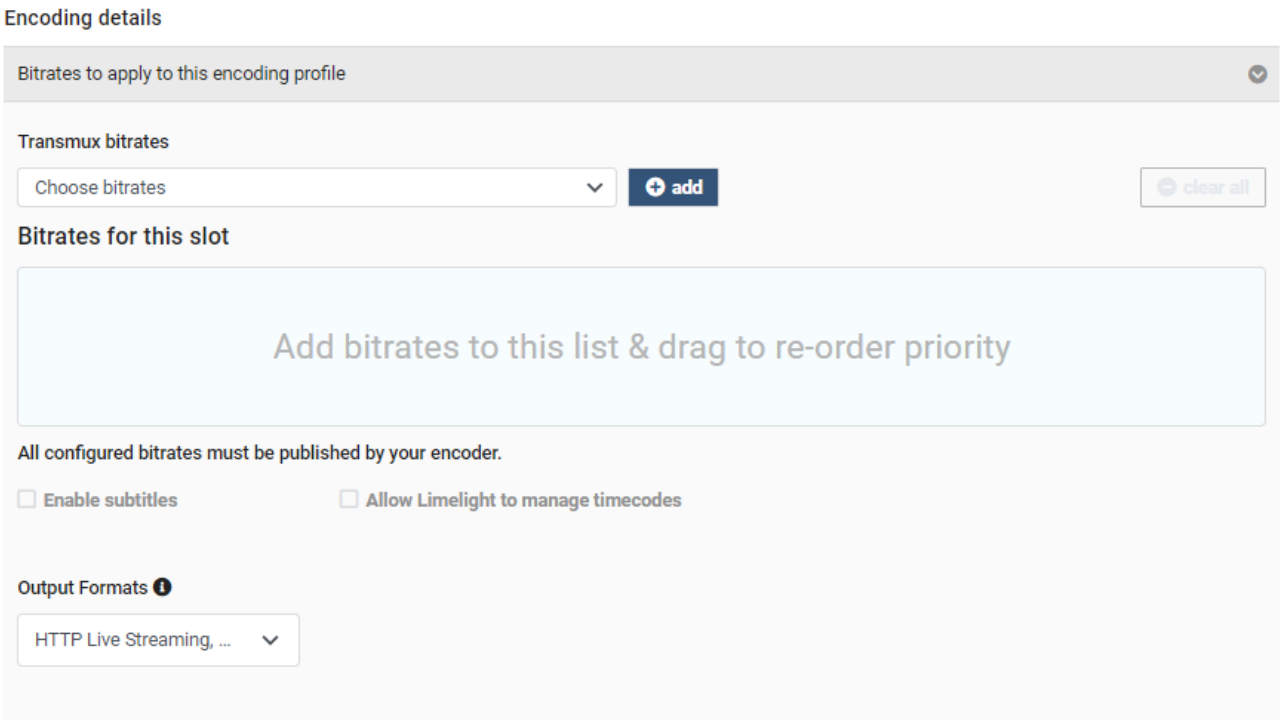

Field Name	Description/Instructions	Notes
Transmux bitrates	The bitrates you want to output from the slot. Choose one or more bitrates from the drop-down menu, then click the **add** button to add the bitrates to the 'Bitrates for this slot' field.	Each selection is an encoding profile Your encoder must publish all configured bitrates.
Bitrates for this slot	A suggested set of video and audio bitrates is available in the drop-down boxes, but you may also enter your custom bitrates in those boxes. If you have 'Bitrate Order' enabled for your account, each bitrate has a drag handle to its left. Drag and drop bitrates to reorder them. The order number determines the order in which the bitrate URL appears in the output manifest file.	You must select at least one bitrate. All configured bitrates must be published by your encoder to MMD Live for the slot to function correctly. Total bitrate (video + audio) is automatically calculated and displayed in the row for each bitrate. The summation of all totals is displayed at the top of the bitrate list. To remove a bitrate, hover over its row and click the **remove** button.
• **Enable subtitles** • **Allow Limelight to**	See Subtitles and Timecodes.	

Field Name	Description/Instructions	Notes
manage timecodes • **Output Formats**		

Realtime Streaming Slots

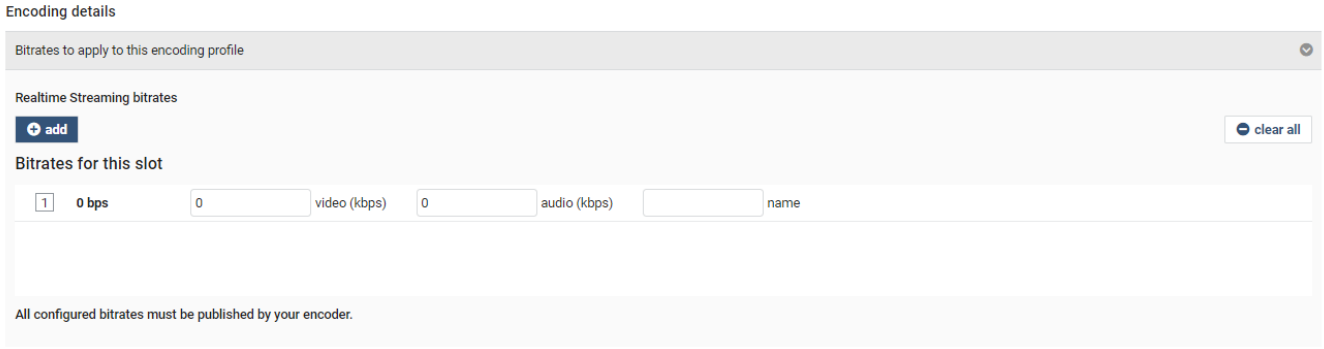

Field Name	Description/Instructions	Notes
Realtime Streaming bitrates	The bitrates you want to output from this slot. By default, none are selected, but you must choose at least one to proceed. You can create bitrate profiles by entering a name for each bitrate. When Adaptive Bitrate is present, the name is displayed in video players instead of the default 480, 720, and so on. The name is also appended to the slot's stream name. 1. Click the **add** button to add one or more bitrates. 2. For each bitrate, configure the video and audio rates. 3. Enter a profile name in the **name** field. If necessary: • Click the **clear all** button to remove all bit rates in the list. • Hover your mouse pointer over a row and click the **remove** button to remove that bitrate.	Each selection is an encoding profile. Your encoder must publish all configured bitrates. When Adaptive Bitrate (ABR) is present in the player, the name is displayed in the player as a stream option instead of the individual Bitrates.
Bitrates for this slot	If you have 'Bitrate Order' enabled for your account, each bitrate has a drag handle to its left. Drag and drop bitrates to reorder	You must configure at least one bitrate. All configured bitrates must be published by your encoder to MMD Live for the slot to function correctly.

Field Name	Description/Instructions	Notes
	them. The order number determines the order in which the bitrate URL appears in the output manifest file.	Total bitrate (video + audio) is automatically calculated and displayed in the row for each bitrate. The summation of all totals is displayed at the top of the bitrate list. To remove a bitrate, hover over its row and click the **remove** button.

Subtitles and Timecodes

Beneath the bitrates list for some slots are checkboxes for managing subtitles and timecodes and a drop-down menu for selecting the output format.

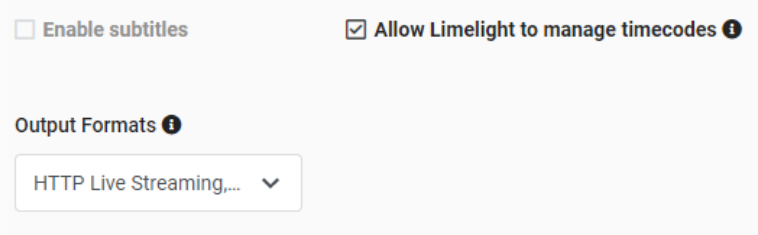

Field Name	Description/Instructions	Notes
Enable subtitles	Whether to enable subtitles in iOS players. Select this option to accurately inform iOS players that subtitles are present in HLS output from MMD Live.	If you don't select the option and the `#EXT-X-MEDIA:TYPE=SUBTITLES` tag is present, but subtitles are not, the iOS player will display a CC menu option to display subtitles even though subtitles are not present.
Allow Limelight to manage timecodes	Whether to allow Limelight to manage time codes or let your encoder manage them. If your encoder does not allow you to enable absolute timecodes in chunks or if you want Limelight to manage timecodes, check this checkbox. If you are recording MMD Live to VoD, this checkbox must be checked.	Absolute timecodes in chunks are necessary to enable a seamless transition from primary to backup ingest in case a primary ingest fails. Timecodes allow the failover mechanism to seamlessly switch to the backup stream at the correct time, resulting in little or no interruption to a viewer's experience when watching the live stream. By default, the checkbox is checked for transmux and transcode slots.
Output Formats	Desired video output standard.	The slot will produce output in the formats that you selected.

Field Name	Description/Instructions	Notes
	Select one or more output formats; then click the 'Apply' entry at the bottom of the drop-down menu.	

Content Security

This section allows you to configure MediaVault and DRM content protection.

Here is a sample configuration for non-Live Push slots:

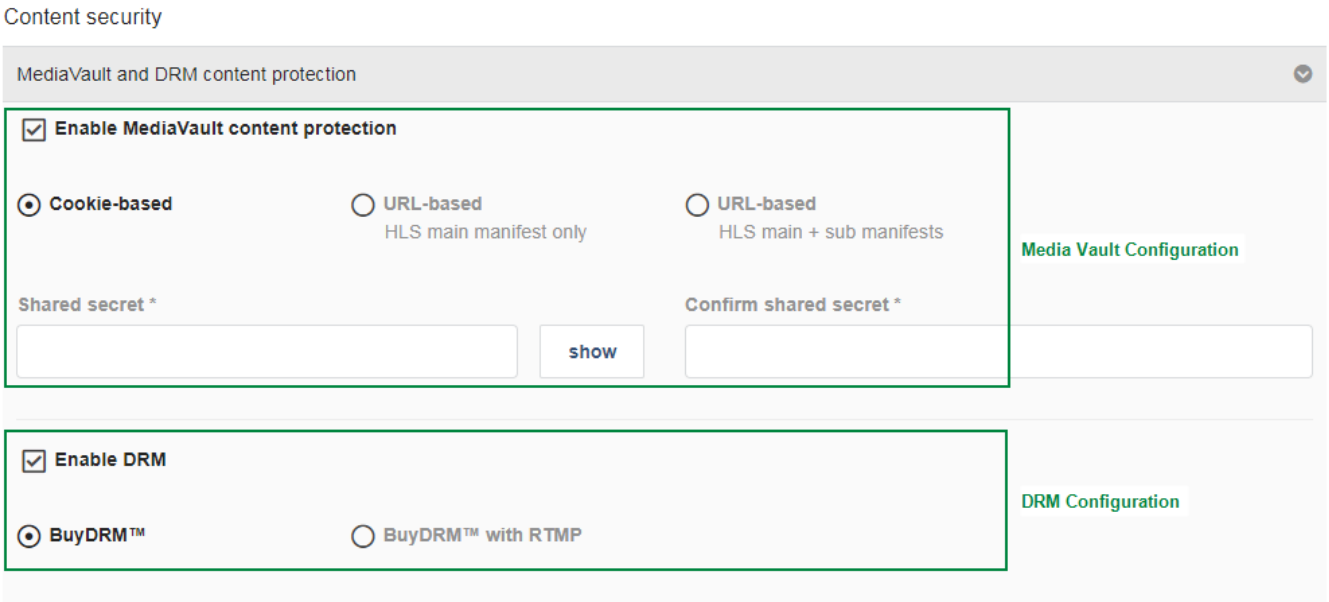

Notes:

For Live Push slots:

- The Enable DRM option is not available.
- The ability to configure **Allow** and **Deny** lists to control viewer access is also available:

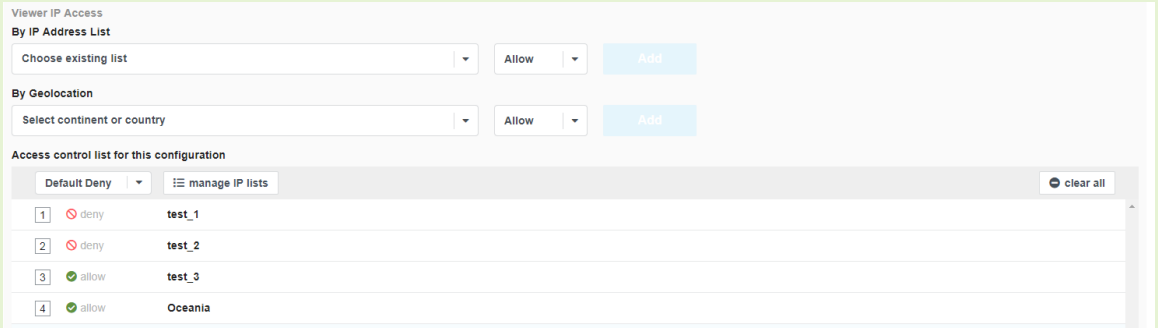

The **Enable Subscribe Validation** option is available for Realtime Streaming slots.

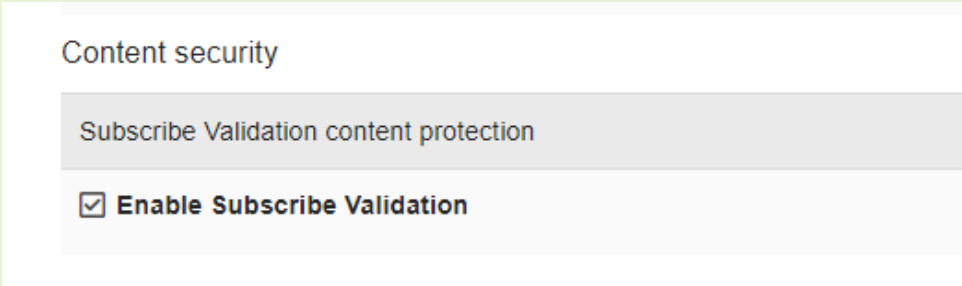

After the slot is created, the Slot Details section provides a link to the location of the related delivery service instance so you can update the service instance and apply access control options. See View Slot Details for additional information.

MediaVault

If you have the MediaVault service option, this section is activated.

Check the **Enable MediaVault content protection** checkbox, then choose the type of MediaVault you would like to implement on this slot:

- URL-based. For URL-based MediaVault, you may choose to MediaVault protect just the main manifest of HLS outputs or both the main manifest and all sub-manifests.
- Cookie-based

You can set your hash secret per slot. You can find more information in the "MediaVault User Guide" or by talking to your Limelight Representative.

> **Note:** Enabling MediaVault causes the Integrated Player Embed Code not to function.

DRM Configuration

If you have the DRM service option, this section is activated.

Check **Enable DRM**, then choose the desired option:

- **BUYDRM**: Enables DRM on MPEG-DASH output; disables all other outputs.
- **BUYDRM_WITH_RTMP**: Enables DRM on RTMP and MPEG-DASH output; disables all other outputs.

Viewer Access Configuration

You can configure **Allow** and **Deny** lists by IP address or geolocation that restrict viewer access to Live Push streams. You can:

- Add existing lists
- Create lists either manually or by uploading a CSV file
- Edit and delete lists

You can also set the action (**Allow** or **Deny**) for any IP address or locations not in the lists that you configured to the stream.

All lists configured for the slot are shown in the **Access control list for this configuration** section at the bottom of the screen:

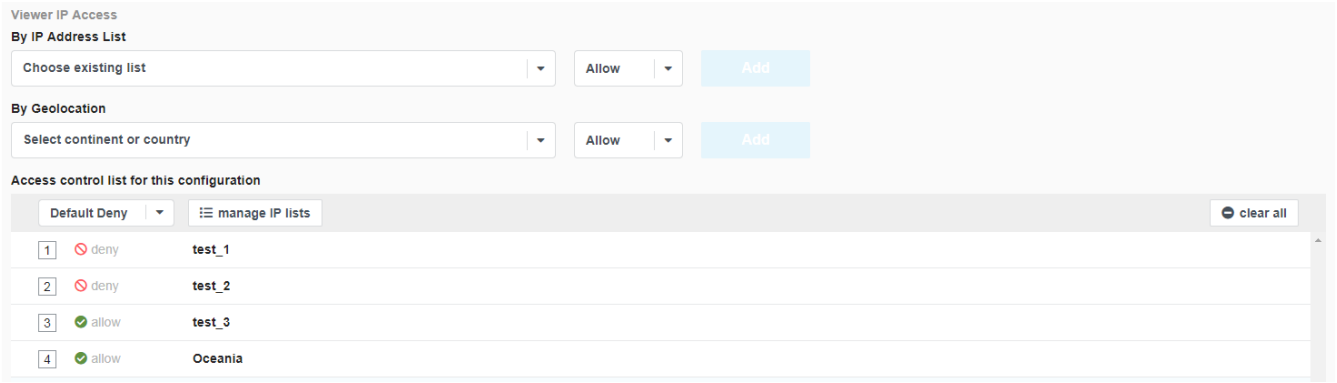

Restricting Access Using Existing Lists

1. Select one or more entries in the **Choose existing** list or **Select continent or country**drop-down menus.
2. Select **Allow** or **Deny** from the drop-down menu, depending on whether you want to allow or deny access to the list.
3. Click the **Add** button to add the Access control list for the slot.

Creating or Cloning an Access List

Click the **manage IP lists** button above the list at the bottom of the page. A dialog with all currently existing lists displays.

1. Choose an action:
   - Create a new list: click the **+ new list** button at the top of the dialog.
   - Clone an existing list: click the **clone** icon on the right side of the list's row.

   The **NEW IP ADDRESS LIST** dialog is displayed.
2. Name the list; then configure IP addresses:
   - Type a new address and press **Enter** to add an address.
   - Click the **x** on an existing address to remove an address.
   - Add to the list by uploading a file of IP addresses. Download the CSV example to get started.
3. Restrict the list to specific accounts by selecting one or more accounts in the drop-down menu.
4. Click the **Save** button.

The new list is now available for you to add, as described in Restricting Access Using Existing Lists.

Viewing Access List Details

Click the **[+]** icon on the left side of the list's name in the dialog.

The entry expands to show read-only details.

Editing an Access List

Begin by clicking the **manage IP lists** button above the list at the bottom of the page. A dialog with all currently existing lists displays.

1. Click the pencil icon on the right side of the list's row in the dialog.
2. Modify any of these fields:
   - Name
   - IP addresses
     - Type a new address or click the **x** on an existing address to remove it.
     - Add to the list by uploading a file of IP addresses. Download the CSV example to get started.
3. Restrict the list to specific accounts by selecting one or more accounts in the drop-down menu.
4. Click the **Save** button.

Deleting an Access List

1. Click the trash can icon on the right side of the list's row in the dialog.
2. Click **Delete** in the confirmation dialog.
3. Control removes the list from the slot's Access control list if it was in the Access control list.

Setting the Stream's Default Access

If, during playback, an IP address or Geolocation that are not in the Access control list try to access the stream, you can set the default access rule (allow or deny access) using the drop-down menu above the slot's Access control list:

Default Allow - automatically allow access from all unknown IP addresses or Geolocations.

Default Deny - automatically block access from all unknown IP addresses or Geolocations.

Removing Selections from the Slot's Access Control List

- Click the **clear all** button at the top right of the list to remove all entries.
- Hover the mouse over an entry and click the **remove** button to delete just that entry.

Clone, Delete, Edit, and View Slots

These sections explain how you can manage your slots:

Clone a Slot

Delete a Slot

Edit a Slot

View Slot Details

View a Slot's Live Stream

Clone a Slot

1. Locate the slot you want to clone and click the **Clone** icon.

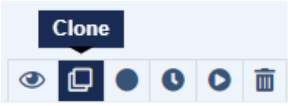

2. Follow the same steps used when configuring a slot. The values are pre-populated with the same information as the original slot except for the slot name. Since the name has to be unique, you must enter a new name.

3. Click the **Submit** button, and the new slot is created.

> **Note:** If you do not change the name or if you enter the name of an existing slot, a warning appears in the right part of the window.

Delete a Slot

> **Note:** You cannot undo or recover a deleted slot. You cannot delete a slot that is in "Pending" state.

1. Locate the slot you wish to delete and click the **Delete** icon.

2. A confirmation dialog is displayed.

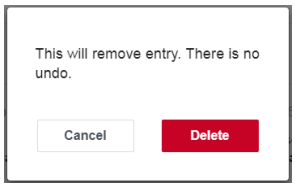

3. Click the **Delete** button.
   The dialog is dismissed and a spinning circle icon is displayed. Upon deletion, the slot is removed from the list.

Edit a Slot

> **Note:** You can edit Live Push slots only.

1. Click the pencil icon on the right side of the list's row in the dialog.
2. Modify any of these fields:
   - Name
   - IP addresses
     - Type a new address or click the **x** on an existing address to remove it.
     - Add to the list by uploading a file of IP addresses. Download the CSV example to get started.
3. Restrict the list to specific accounts by selecting one or more accounts in the drop-down menu.
4. Click the **Save** button.

View Slot Details

> **Note:** For information about viewing Live Push slots details, please see the Video Live Push Guide.

1. Locate the slot and click the **Details** icon (or click the slot's row).

   - You will see the Ingest URLs and Stream Name to enter into your encoder, and you will see the Playback URLs to use in your player. This example is for a Transcode slot; other slot types are similar:

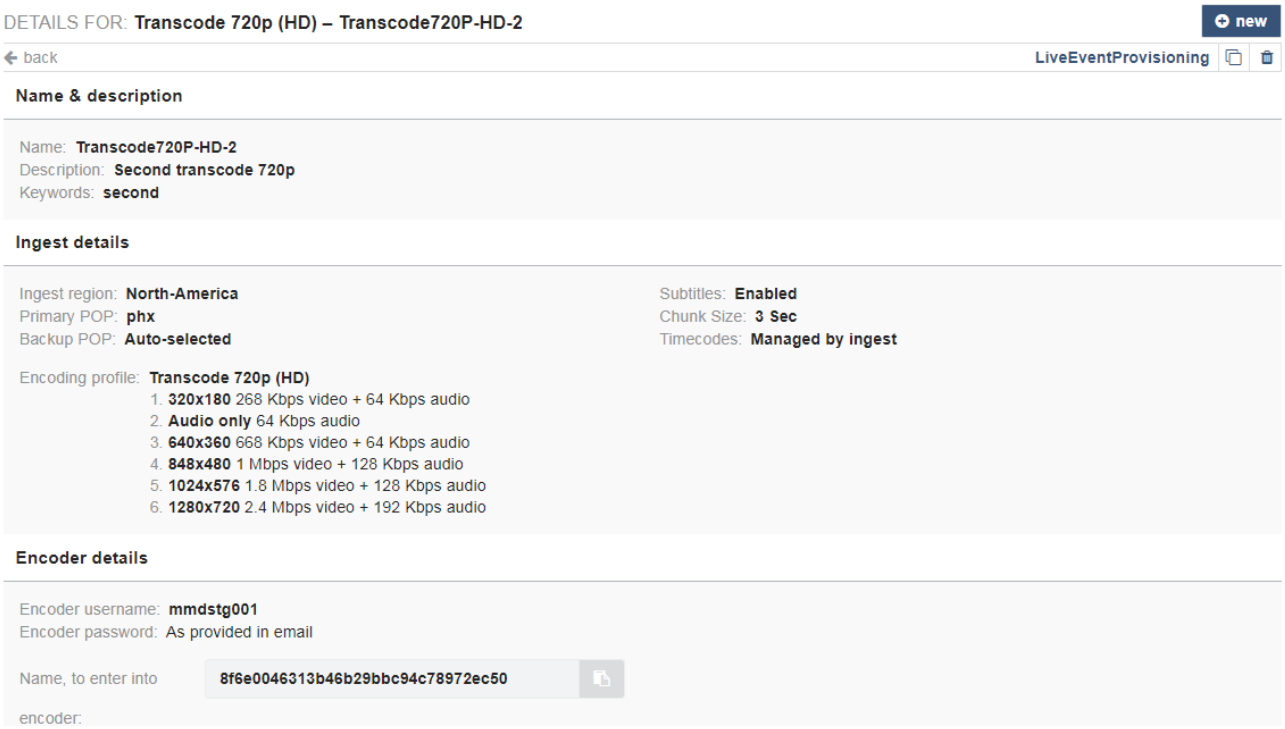

- You can also view the slot's playback URLs. This example is for a Transcode slot; other slot types are similar:

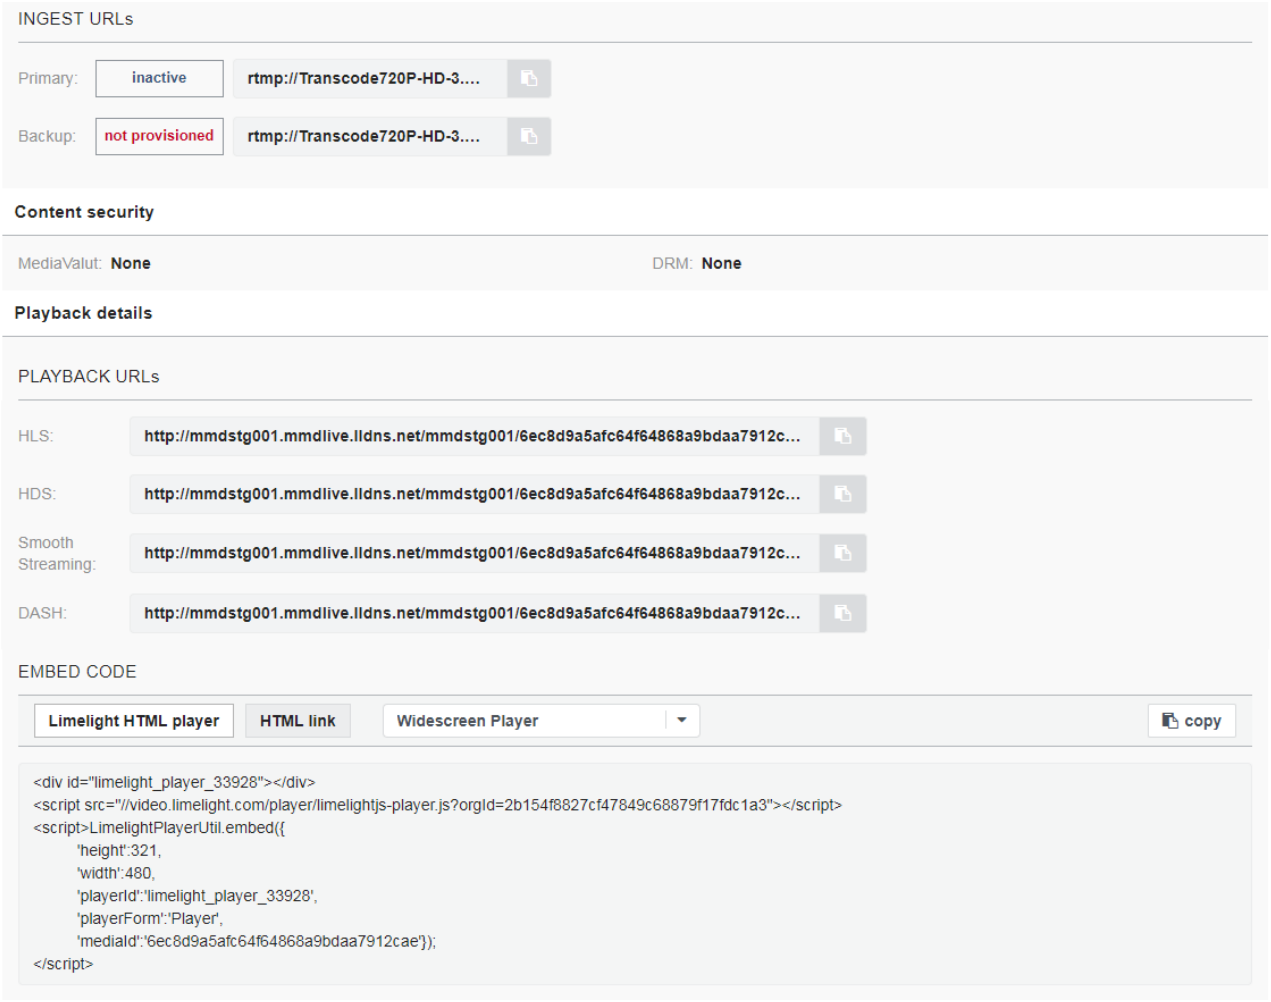

2. If you chose the **Enable Subscribe Validation** option while configuring Content Security for a Realtime Streaming slot, a Delivery service instance was created in which you can configure access control options. To do so, click the **Configure Validation** button in the *Content Security* section.

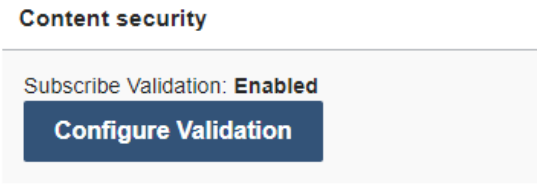

   The service instance opens in Caching & Delivery (v2). See 'Content Security' in Configuring Content Delivery (v2) for further instructions.
3. If you choose to, you may use the Limelight SmartEmbed Player found in the section labeled “EMBED CODE.” Choose an embed code option and player type, then copy the embed code and place it on your website. The embed code will play your live stream. For more information about SmartEmbed, see the Player Embedding Guide.

> **Note:** If you enable MediaVault, the EMBED CODE section is not displayed because MediaVault and embed code are not compatible

4. The Copy icon next to a field allows you to copy the field to the clipboard easily:

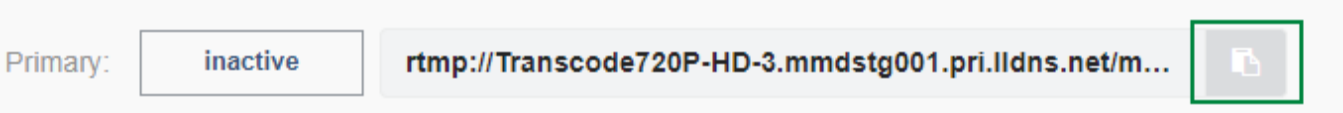

When you click the icon, the browser displays a confirmation that the data has been copied.

5. You can quickly determine if you are streaming to MMD Live ingest servers by looking at the URL status indicator to the right of the Primary and Backup **Ingest URLs**:
   - "Active" means the encoder is currently publishing to that ingest.
   - "Inactive" means the encoder is not currently publishing to that ingest.
   - "Error" or "Not Provisioned" means there was an error when querying the ingest for status. Please contact Customer Support.

View a Slot's Live Stream

> **Note:** You cannot view a Live Push slot's stream.

To view a slot's live stream:

1. Locate the desired slot and click the **Player** icon.

2. A new tab opens in your browser with the HTML player embedded in it. The player plays the live stream.

> **Notes:**
> The **Player** icon is only enabled if MediaVault is not enabled for the slot. See MediaVault for additional information.
>
> The embedded player look and feel is defined by the options configured for the player in the EMBED CODE section of the slot details screen. See View Slot Details for additional information.

Using Your Slot

> **Note:** For instructions about using Live Push slots, see 'Use Your Slot' in the Video Live Push Guide.

Once your slot is configured, you can begin streaming to it right away. The slot is available to you whenever you want to use it.

Start by setting up your encoder with the information provided in the “Slot Details” screen. See 'Setting Up Your Encoder' in the MMD Live Streaming Guide for encoder requirements and recommended settings.

Be sure to use the specifications of your slot type when setting up your encoder. It's also important to publish to both primary and backup publishing URLs and use an absolute Timecode in your encoder to provide maximum failover protection for your live stream.

Once you are streaming, use one of the playback URLs shown in the “Slot Details” screen in your video player or app. Or you can use the Limelight SmartEmbed on your website or blog, which loads a player that automatically contains your playback URL.

Using Secure Playback URLs

MMD Live and Live Push support the delivery of live streams over secure URLs. Simply change your playback protocol from `http://` to `https://` to take advantage of secure delivery.

Manage

My Account

The 'My Account' section of *Limelight Control* lets you update your contact information and password. The section also lets you manage any automatically-recurring report emails and alerts you've created.

To view the *Edit My Account* page, select *My Account* in the **Welcome** drop-down menu.

Editing Your Profile

Your profile information is in the ***Profile*** tab, which is visible by default when you access the *Edit My Account* page.

Contact Information

You can change your first name, last name and email address at any time in the fields provided.

> **Note:** Changing your email address does not change your user name. *Limelight Control* user names are permanent.

Locale & Timezone

To change the language displayed by the *Limelight Control* user interface, select a language from the **Locale**drop-down menu.

To select the default timezone for report data, choose a GMT timezone from the **Time Zone**drop-down menu.

Default Account

If you have more than one Account (also known as a "shortname"), you can use the **Default Shortname** drop-down menu to choose the company and account you want to be selected when you log in.

> **Note:** You can change the selected Account at any time in the Company/Account drop-down menu in the top ribbon.

Default Landing Page

When you first log in to *Limelight Control*, the Dashboard is automatically displayed, but you can configure another page to be the default. Click the **Landing Page** drop-down menu and select the desired page.

Saving

When you've finished with your changes, click the **Save** button.

> **Hint:** You may need to log out and then log back in again for the new settings to take effect.

Changing Your Password

To change your password:

1. Click **Change Password.**
2. Enter and confirm the new password in the fields provided.
3. Click the **Save** button.

> **Note:** The last five passwords may not be entered as a new password. Additional details on acceptable password length and content are displayed in Limelight Control when a new password is entered.

API Shared Key

You can also view your API key or create a new one. The API key is needed if you want to access any of the APIs associated with *Limelight Control*.

To view your API key, click **Show my Shared Key**. To create a new API key, click **Generate a new Shared Key**.

Managing Recurring Report Emails

The **Recurring Emails** tab allows you to view, edit, and delete any recurring emails that you configured in the Traffic Report and the Status Codes Report.

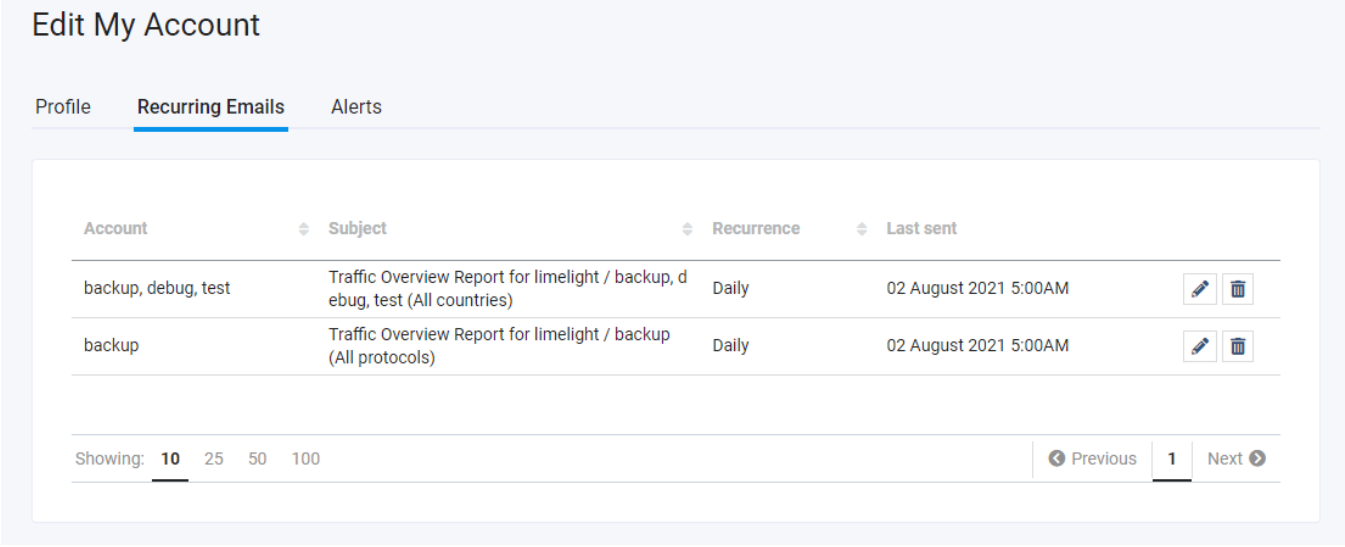

The tab contains a list of configurations, each with the fields in the following table.

Field	Description
Report	Report in which the email was configured.
Account	Associated Limelight Account.
Subject	Email subject as seen by recipients.
Recurrence	How often the email is transmitted.
Last Sent	Most recent date and time that the report was transmitted.

On the right side of each row are icons for editing and deleting the configuration.

The tab contains pagination controls you can use to browse current configurations.

Editing a Recurring Report Email

To edit a recurring report email:

1. Click the **edit** icon.
   The RECURRING EMAIL dialog is displayed. With a few differences, the selections in dialog's fields are the same whether the email was generated from the Traffic Report or the Status Codes Report.
2. Edit the fields using information in the Fields in the *RECURRING EMAIL* Dialog.
3. Click the **Save** button.

Deleting a Recurring Report Email

To delete a recurring report email, click the **delete** icon.

Managing Alerts

The ***Alerts*** tab allows you to view, edit, and delete alerts that you created in the Traffic Report and the Status Codes Report.

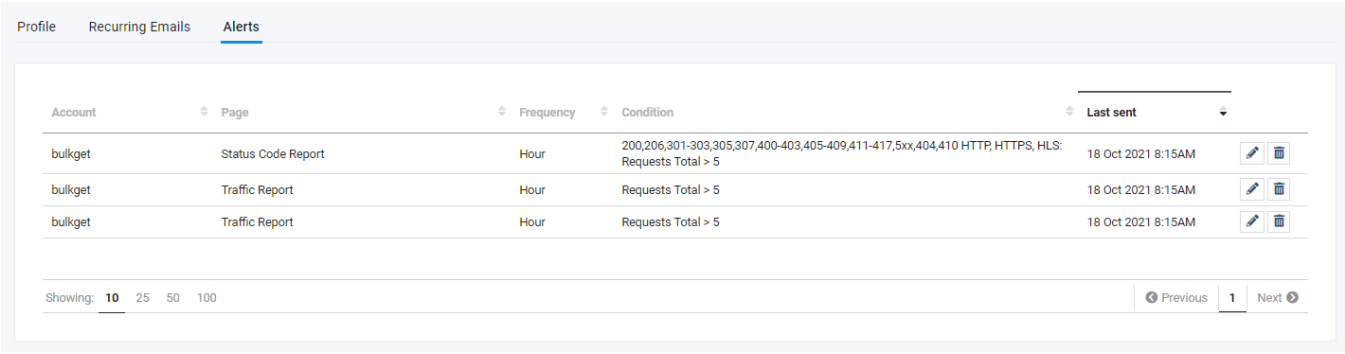

The tab contains a list of configurations, each with the fields in the following table.

Field	Description
Account	Associated Limelight Account.
Page	Report in which the email was configured.
Frequency	How often the email is sent.
Condition	The circumstances to trigger the email.
Last Sent	Most recent date and time that the report was transmitted.

Each entry has an **edit** and **delete** icon on the right side of its row.

Editing an Alert

1. Click the **edit** icon for the alert you wish to modify.
   The REPORT ALERT dialog is displayed.
2. Make changes in the dialog. Not all fields are editable. See Editable Fields in the REPORT ALERT Dialog for details.
3. Click the **Save** button.

Editable Fields in the REPORT ALERT Dialog

Field	Description / Instructions
Recipient	Primary notification recipient. Enter a single primary recipient email.
CC	Stands for "carbon copy." Additional email recipients. Enter one email or multiple emails separated by commas.
Subject	Email subject. Enter the text for the email's subject line.
Message	Email body text. Enter the text for the email body.
Show UI notification	If this field is checked, then when the threshold has been crossed, Limelight Control displays an alert popup in the user interface in addition to sending an email.

Deleting an Alert

1. Click the **delete** icon for the alert you wish to delete.
2. Confirm that you want to delete the alert in the dialog that displays.

SmartPurge

SmartPurge is Limelight's innovative system for removing content from CDNcache.

Why Purge?

Objects are normally updated in or removed from cache during “freshness checks” with your origin. For a given object, a freshness check is initiated when a request has been made for the object, and the object’s TTL (Time To Live) has expired.

In general, setting object TTL is the best and most efficient way to manage cached content. For example, a news site may need to provide rapid updates to a breaking video story. The video can be updated in cache as quickly as desired by assigning it a low TTL value using an HTTP response header. In most cases, there is no need to remove the video from cache directly.

However, there are special cases where content needs to updated on the next user request or even pro-actively removed from cache as soon as possible. This is known as "purging the cache" or just "purging". Examples of when purging might be necessary include:

- Text is misspelled in the caption of a newly-uploaded video, and you need to update the video in cache as quickly as possible.
- You discover that some of your cached content is infringing a copyright and need to delete the content from cache as soon as possible.
- You lose a contract with a content provider and are obligated to delete the provider’s content from your cache as soon as possible.
- During a full website update, when you need to quickly update many related website objects (images, text, video, etc.) at the same time.

Limelight's *SmartPurge* executes purge operations more quickly and reliably than older technologies. The advanced version of *SmartPurge*, *SmartPurge Plus*, provides additional features, including higher purge queue priority, and additional API features such as unlimited callbacks.

You can access *SmartPurge* through either the Limelight Control Portal or the SmartPurge REST API. For more on purging, please see the Limelight SmartPurge Data Sheet and Intelligent High-Speed Purging.

Notes:
Access to SmartPurge is granted by default for Company Admins but must be explicitly granted for uses with other roles.

SmartPurge also purges "negatively cached" content, such as HTTP error responses (404, etc.). This lets you remove negatively cached content when the content lifetime has been extended, such as when the origin includes a `Cache-Control` header that has a `max-age` value much greater than the default.

Purging MMD OD assets is supported only by public manifest on the Enter URLs Tab.

SmartPurge Page Overview

The SmartPurge page has two tabs: Requests and Templates.

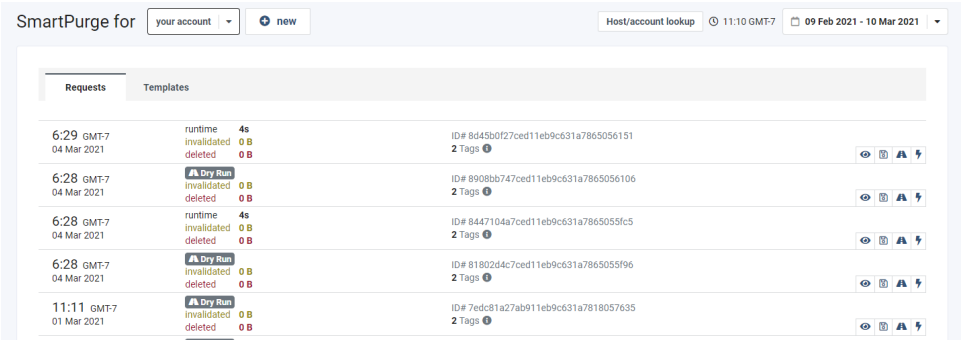

Controls at the top of the page allow you to choose an account and create a new purge request, create a new template, or create a new request from a template.

From the **Requests** tab only, you can filter the list of requests by date range and do a host/account lookup.

Requests Tab

The **Requests** tab lists all currently configured purge requests. Each request shows this information:

- Time, timezone, and purge date.
- A 'Dry Run' icon if the request was a dry run.
- Total bytes invalidated and deleted.
- Request ID.
- Number of patterns, URLs, and tags in the request.
- Information icon beside patterns, URLs, and tags. Click the icon to view details.

What You Can Do on the Requests Tab

Use the icons on the right of each row to :

- View a Request's Stats (Results)
- Save a Request as a Template
- Do a Dry Run
- Rerun the request. See Doing a Purge.

Templates Tab

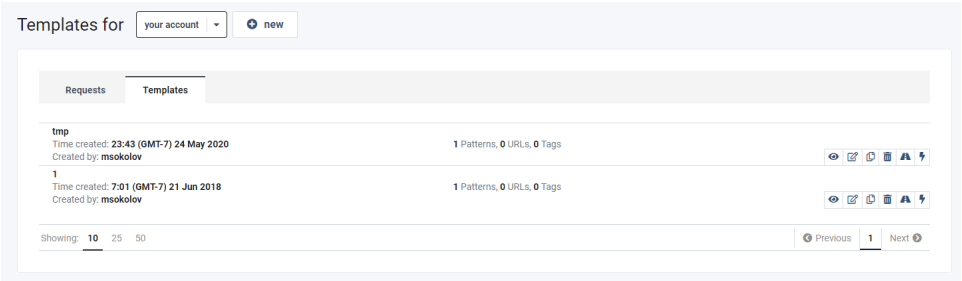

The **Templates** tab lists purge requests that have been saved as templates for easy reuse and minimal duplication of effort. Each template shows this information:

- Template name.
- Date and time the template was created.
- User that created the template.
- Number of patterns, URLs, and tags in the template.

What You Can Do on the Templates Tab

Use the icons on the right of each row to:

- View Template Summary
- Edit a Template
- Duplicate a Template
- Delete a Template
- Do a Dry Run
- Do a Purge

Creating a New Purge Request

You can create a purge request from an existing template or from scratch using the *New Purge request for* page.

1. Click the **+ new** button and select **Request** from the subsequent drop-down menu.
2. Choose an option for creating the request:
   - **Request**: See Creating a Purge Request from Scratch.
   - **From existing template**: See Creating a Request from a Template.

Creating a Purge Request from Scratch

SmartPurge provides four tabs for specifying objects to purge. You can use all or any of the tabs for one request. Change you make in one tab are reflected in the others. The tabs are:

- Build Patterns
- Enter URLs
- Apply tags
- Upload file

Build Patterns Tab

This tab allows you to configure items to purge by specifying patterns. Patterns can include any portion of the Origin URL, path, filename, or extension.

Setting	Information Requested	Purpose	Selecting the Right Option
Protocol	Whether to purge objects that were previously requested with the HTTP protocol, the HTTPS protocol, or both	The protocol is part of the Cache Key for cached objects	In the **Protocol**drop-down menu, select the protocol(s) used in your published links to the content you want to purge.
Published Host	Which of your existing Published Hosts should be used when matching cached objects	The **Published Host** is part of the Cache Key used to retrieve cached objects	In the **Published Host**drop-down menu, select the desired entry.
Published Path	Whether to limit matches to a specific URL path when matching cached objects	Select this option to limit	If desired, enter a path in the **Published Path**

Setting	Information Requested	Purpose	Selecting the Right Option
		the purge operation to content in a specific folder or folders.	field. **Note:** Pathnames are case sensitive
Include query string	Whether to include the query string, if any, when matching cached objects	Select this option to limit the purge operation to URLs with a specific pattern in the query string	If desired, check the **Include query string** checkbox to include cached query strings in the match **Note:** Query strings are case sensitive
Origin URL / Cache Key	Populated after you click the **Apply** button. See SmartPurge Pattern Details for more information.		

SmartPurge Pattern Details

SmartPurge patterns can include any portion of the Origin URL, path, filename, or extension. For example, `*.mysite.com` will match all content from Published Hostnames associated with mysite.com, `/abc/` will match all directories named “abc” across all hostnames, and `*.mp4` will match all MP4 files across all hostnames.

Notes:

In the **Origin URL or Pattern** field,

- **Case Sensitivity**. Patterns are case sensitive.
- **Wildcards**. The wildcard character (“*”) is allowed but cannot be used by itself.
- The use of (“*”) alone is only supported in the SmartPurge API.
- **Spaces & Other Normally Encoded Characters**. Spaces in patterns must be URL encoded (replaced with "%20"), and other characters normally URL encoded must be similarly replaced.
- **Folder Paths**. For patterns that begin with a folder name (including root directory folders), you must prefix the folder name with "*/". This is because *SmartPurge* compares the pattern to the entire cache key, which begins with the protocol and Origin URL, not just the folder path,

If you are having trouble creating a pattern, click the **Host/account lookup** button at the top right of the tab headers. Entering an exact Published Host will yield the corresponding exact Origin URL, including URL encoded characters. You can then use the exact Origin URL as the basis for your pattern.

When you've finished entering the above settings, click the **Apply** button, and the Cache Key prefix(es) (**Protocol** + **Published Host** + **Published Path**) will be shown in the *Origin URL / Cache Key* section.

You can now continue entering purge settings:

Setting	Information Requested	Purpose	Selecting the Right Option
What do you want to purge?	Which pattern type(s) to use when matching cached objects	At least one type of pattern must be specified before *SmartPurge* can begin matching cached objects	Click one or more of these pattern types: • **Exact origin URLs** - if you know exactly which Origin URLs you want to purge. You can also load them from a file using the **Upload file** tab. • **File extensions** - to purge all files with the file extensions you enter in the associated field • **File names** - to purge all files with the file names you enter in the associated field • **Directories** - to purge all files in the directories, you select **Directories**. To include all files below the selected directory, choose **Include subdirectories & contents**. • **Entire sites/origin hosts** - to purge all files from the paths shown in **Origin URL / Cache Key**, without limitation. **Note:** This option is mutually exclusive with **Exact origin URLs** - only one of the two may be selected. For File extensions, File names, and Directories, you must enter values in the text box, and press the Enter key after entering each value. You can remove an entry by clicking the x on the right side of the entry. For each type of match you select, you must: • Select **Delete** or **Invalidate** to the right of the text box. • Click the **Add** button to save each entry, after which it will appear in the **List of patterns/URLs/tags** field • Specify whether to Invalidate or Delete the matched objects **Notes:** • Patterns are case sensitive • A maximum of 100 URLs and 100 Patterns may be submitted per request, and a maximum of 60 URLs and 60 Patterns may be submitted per minute per Account.

Setting	Information Requested	Purpose	Selecting the Right Option
List of pat-terms/URLs/ tags	Whether you want to delete or invalidate an entry in the list. Whether you want to delete an entry from the list.	This section contains a list of the patterns, URLs, and tags to purge that you con-figured in all four tabs. You can use the section to make a final decision whether: • To delete or invalidate the objects specified by the entries. • You indeed want to do a purge by removing one or all entries from the list.	To change the purge type, click **Delete** or **Invalidate**. Remove single items as needed from the section by clicking the trash can icon on the right side of the row. Remove all items by clicking the **clear all** button above the list on the right side of the screen.
Notes	Optional notes that you can refer to later when browsing historical con-figuration changes	Notes allow you to include additional information for others (what files were targeted, why the purge was needed, etc.)	If you want to save notes with your purge request, just enter them in the **Notes** field
Save this request as a template	Whether you want to save the purge request for future re-use	If you plan to purge the same objects again in the future, you may want to re-use the purge request	To save the purge request as a template, check the **Save this request as a template** checkbox. **Note:** You can also save the request as a template, without sending a purge request, by clicking **Save** (instead of **Purge**) when you are finished.
Notify when completed	Whether you want to be notified via email when your purge request has been completed	A notification is sent when the purge is complete. This provides the option to take additional action or inform others when the purge is complete	To be notified via email when your purge request has completed, check the **Notify when completed** checkbox, and enter the email address(es) to notify. **Note:** In the notification email, the purged object count is approx-imate and may increase after the email has been sent

Note:
Information in these fields are reflected in all tabs:

- List of patterns/URLs/tags
- Notes
- Notify when completed and email addresses

When you have finished making changes, click **Save as template** to save your request so you can continue work later, **Cancel** to discard your settings, **Dry run** to test your settings, or **Purge** to submit a purge request with your settings.

Note: When you click **Save as template**, the *Create template for page* displays so you can save the request as a template.

See:

Saving a Request as a Template

Doing a Purge

Doing a Dry Run

Enter URLs Tab

This tab allows you to specify exact URLs to purge.

Setting	Information Requested	Purpose	Selecting the Right Option
Protocol	Whether to purge objects that were previously requested with the HTTP protocol, the HTTPS protocol, or both	The protocol is part of the Cache Key for cached objects	In the **Protocol**drop-down menu, select the protocol(s) used in your published links to the content you want to purge
Exact Published URL	The exact *Published URL* to purge	Select this option to purge specific *Published URLs* rather than using a pattern match	Enter the exact *Published URL* in the field. **Notes:** • Asterisk (*) characters are accepted, but are interpreted as literals, not wildcards. • URLs are case sensitive
Invalidate/Delete	Whether to Invalidate or Delete the matched objects	In many cases, it is acceptable to flag objects for replacement on the next user request. However, in some cases, you	To update matching objects on the next user request, choose **Invalidate** under Purge Type. To proactively remove objects from the cache, choose **Delete**. When you are finished, click the **Add** button, and the target URL you specified will appear in the **List of patterns/URL/stags** field.

Setting	Information Requested	Purpose	Selecting the Right Option
		may need to remove content from cache pro-actively.	
List of pat-terms/URLs/ tags	Whether you want to delete or invalidate an entry in the list. Whether you want to delete an entry from the list.	This section contains a list of the patterns, URLs, and tags to purge that you configured in all four tabs. You can use the section to make a final decision whether: • To delete or invalidate the objects specified by the entries. • You indeed want to do a purge by removing one or all entries from the list.	To change the purge type, click **Delete** or **Invalidate**. Remove single items as needed from the section by clicking the trash can icon on the right side of the row. Remove all items by clicking the **clear all** button above the list on the right side of the screen.
Notes	Optional notes that you can refer to later when browsing historical con-figuration changes	Notes allow you to include additional information for others (what files were targeted, why the purge was needed, etc.)	If you want to save notes with your purge request, just enter them in the **Notes** field
Notify when completed	Whether you want to be notified via email when your purge request has been completed	A notification is sent when the purge is complete. This provides the option to take additional action or inform others when the purge is complete	To be notified via email when your purge request has completed, check the **Notify when completed** checkbox, and enter the email address(es) to notify

Note:
Information in these fields are reflected in all tabs:

- List of patterns/URLs/tags
- Notes
- Notify when completed and email addresses

Note: A maximum of 100 URLs and/or 100 Patterns may be submitted per request, and a maximum of 60 URLs and/or 60 Patterns may be submitted per minute (i.e., one URL or Pattern per second) per Account (shortname). For example: after submitting a purge request with 100 URLs, it is necessary to wait for 100 seconds before submitting the next request.

When you have finished making changes, click **Save as template** to save your request so you can continue work later, **Cancel** to discard your settings, **Dry run** to test your settings, or **Purge** to submit a purge request with your settings.

Note: When you click **Save as template**, the *Create template for page* displays so you can save the request as a template.

See:

Saving a Request as a Template

Doing a Purge

Doing a Dry Run

Apply Tags Tab

This tab allows you to purge by a metadata tag that was previously supplied with the cached object via headers.

Setting	Information Requested	Purpose	Selecting the Right Option
Purge by tag	One or more tags on which to base the purge.	At least one tag must be specified before *SmartPurge* can begin matching cached objects	Enter one or more tags. Press enter after you type each. When you have entered all tags, click **Add** to save the tags in the **List of patterns/URLs/tags** field. **Notes:** • You can also enter patterns and URLs. • You can paste tags into the field. • If you make an invalid entry, tag specifications and a warning are displayed above and below the field.
	Whether to invalidate the items or remove them.	This setting allows	Select **Delete** to completely

Setting	Information Requested	Purpose	Selecting the Right Option
		you to intelligently manage your cache and minimize requests to the origin.	remove the item from cache. This has a performance impact because the next time the item is requested, the item has to be retrieved from your origin. Select **Invalidate** to simply hide objects from user access and cause an If-Modified-Since request against your origin, and will only cause cache-fill for objects that have changed.
List of pat-terms/URLs/ tags	Whether you want to delete or invalidate an entry in the list. Whether you want to delete an entry from the list.	This section contains a list of the patterns, URLs, and tags to purge that you configured in all four tabs. You can use the section to make a final decision whether: • To delete or inval-idate the objects specified by the entries. • You indeed want to do a purge by removing one or all entries from the list.	To change the purge type, click **Delete** or **Invalidate**. Remove single items as needed from the section by clicking the trash can icon on the right side of the row. Remove all items by clicking the **clear all** button above the list on the right side of the screen.
Notes	Optional notes that you can refer to later when browsing historical configuration changes	Notes allow you to include additional information for others (what files were targeted, why the purge was needed, etc.)	If you want to save notes with your purge request, just enter them in the **Notes** field
Notify when completed	Whether you want to be notified via email when your purge request has been completed	A notification is sent when the purge is complete. This provides the option to take additional action or inform others when the purge is complete	To be notified via email when your purge request has completed, check the **Notify when completed** checkbox, and enter the email address(es) to notify

Note:
Information in these fields are reflected in all tabs:

- List of patterns/URLs/tags
- Notes
- Notify when completed and email addresses

When you have finished making changes, click **Save as template** to save your request so you can continue work later, **Cancel** to discard your settings, **Dry run** to test your settings, or **Purge** to submit a purge request with your settings.

Note: When you click **Save as template**, the *Create template for page* displays so you can save the request as a template.

Upload file Tab

This tab allows you to upload a CSV file that contains specifications of objects to purge.

Setting	Information Requested	Purpose	Selecting the Right Option
Select a file	The file containing specifications of objects to purge.	At least one specification must be in the file before *SmartPurge* can begin matching cached objects	Click in the associated field and select a file of purge specifications to upload. If the file contains errors, a description of the first error is displayed below the field. Fix the error and upload the file again. You can download a sample file by clicking the link under the field. When the upload is complete, the patterns in the file will appear in the **List of patterns/URLs** field. See Purge File Format Details for additional information.
List of patterms/URLs/tags	Whether you want to delete or invalidate an entry in the list. Whether you want to delete an entry from the list.	This section contains a list of the patterns, URLs, and tags to purge that you configured in all four tabs. You can use the section to make a final decision whether: • To delete or invalidate the objects specified by the entries. • You indeed want to do a purge by removing one or all entries from the list.	To change the purge type, click **Delete** or **Invalidate**. Remove single items as needed from the section by clicking the trash can icon on the right side of the row. Remove all items by clicking the **clear all** button above the list on the right side of the screen.

Setting	Information Requested	Purpose	Selecting the Right Option
Notes	Optional notes that you can refer to later when browsing historical configuration changes	Notes allow you to include additional information for others (what files were targeted, why the purge was needed, etc.)	If you want to save notes with your purge request, just enter them in the **Notes** field
Save this request as a template	Whether you want to save the purge request for future re-use	If you plan to purge the same objects again in the future, you may want to re-use the purge request	To save the purge request as a template, check the **Save this request as a template** checkbox. **Note:** You can also save the request as a template, without sending a purge request, by clicking **Save** (instead of **Purge**) when you are finished.
Notify when completed	Whether you want to be notified via email when your purge request has been completed	A notification is sent when the purge is complete. This provides the option to take additional action or inform others when the purge is complete	To be notified via email when your purge request has completed, check the **Notify when completed** checkbox, and enter the email address(es) to notify

Note:
Information in these fields are reflected in all tabs:

- List of patterns/URLs/tags
- Notes
- Notify when completed and email addresses

When you have finished making changes, click **Save as template** to save your request so you can continue work later, **Cancel** to discard your settings, **Dry run** to test your settings, or **Purge** to submit a purge request with your settings.

Note: When you click **Save as template**, the *Create template for page* displays so you can save the request as a template.

See:

Saving a Request as a Template

Doing a Purge

Doing a Dry Run

Purge File Format Details

Each line within the text file is case sensitive and should be in the following format:

```
pattern,exact_match,evict,include_query_string,tag
```

Parameter	Description	Syntax
`accountname`	The LimelightAccount (shortname) to purge.	string
`pattern`	A pattern that describes the content to purge. Matched against content in all *Origin URLs* associated with the Account. **Note:** The pattern is matched against Origin URLs only (not Published URLs).	string • a *partial URL*, or • a *pattern*, or • a *fully qualified URL*
`exact_match`	Whether to treat the `pattern` as a partial URL or pattern, or as an exact match for a fully qualified URL **Note:** exact_match patterns can operate on Published URLs only.	integer (`0` or `1`) • `0` - *partial URL* or *pattern* • `1` - *fully qualified URL*
`evict`	Whether to *Invalidate* or *Evict* targeted content. Invalidating ensures that freshness checks are made for matching objects on the next user request. Evicting causes them to be deleted from cache entirely.	integer (`0` or `1`) • `0` - *Invalidate* • `1` - *Evict*
`include_query_string`	Whether to include the query string, if any, when matching cached objects	integer (`0` or `1`) • `0` - *Ignore* • `1` - *Include*
`tag`	Tag to invalidate or delete, preceded by integers indicating whether to invalidate or evict.	string

For example, to perform this purge operation:

- Pattern: `http://www.example.com/home/index/` (a partial URL)
- Purge Type: `Invalidate`

The corresponding line in the URL file would be:

```
http://www.example.com/home/index/,0,0,0
```

As another example to evict objects identified by `tag1`, and invalidate objects identified by tag2, the corresponding lines would be:

```
,0,1,0,tag1
,0,0,0,tag2
```

Creating a Request from a Template

1. Click the **+ new** button, then select **Request from template**.
   The *New purge request from template for* page displays.
2. Select one or more previously saved templates that will form the purge request (up to the 100-Pattern limit)

> **Note:** If you combine multiple templates, notes, and notifications associated with the individual templates are not included in the combined template.

When you have finished making changes, click **Purge** to submit a purge request with your settings, **Dry run** to test your settings, or **Cancel** to discard your settings.

See:

Do a Purge

Do a Dry Run

Creating a New Template

To create a new template, click the **+ new** button and select **Template** from the subsequent drop-down menu. You then can specify what to purge using one of the tabs in the *Create new template for* page.

The *Create new template for* page is identical to the *New purge request for* page.

See Creating New Purge Requests for instructions.

Doing a Dry Run

Starting Tab or Page	Instructions
Requests Tab Templates Tab	1. Click the request's or template's **dry run** icon. 2. In the subsequent dialog, confirm that you want to run the request or template. • Request initiated from **Requests** tab: The request is added to the top of the list in the **Requests** tab. When the request is finished, you can click the new row to view the request's Stats for request page. • Request initiated from the Templates tab: when the request is finished, the Stats for request page displays.
New purge request from template for Page New purge request for Page Stats for request Page	1. Click the **Dry run** button or **dry run** icon, depending on the tab or page. 2. In the subsequent dialog, confirm that you want to run the request. The request is added to the **Requests** tab.

Starting Tab or Page	Instructions
	3. Click the new row to view the request's Stats for request page.
Template Summary Page	1. Click the **dry run** icon. 2. In the subsequent dialog, confirm that you want to run the request. The Stats for request dry run page displays.

Doing a Purge

Starting Tab or Page	Instructions
Requests Tab Templates Tab	1. Initiate the request. • From the Requests Tab: click the request's **rerun** icon. • From the Templates Tab: click the template's **purge** icon. 2. In the subsequent dialog, confirm that you want to run the request or template. The request is added to the **Requests** tab. 3. Click the new row to view the request's Stats for request page.
New purge request for Page New purge request from template for Page Stats for request Page	1. Click the **Purge** button or **rerun** icon depending on the page. 2. In the subsequent dialog, confirm that you want to run the request. The request is added to the **Requests** tab. 3. Click the new row to view the request's Stats for request page.

Other Request Tasks

Viewing a Request's Stats (Results)

1. Locate the desired request in the **Requests** tab.
2. Click the request or click the request's **view stats** icon.
   The Stats for request page page displays.

Saving a Request as a Template

Starting Tab or Page	Instructions
New purge request for Page Requests Tab Stats for request Page	1. Initiate the save action. • From 'New purge request for' page: a. Click the **Save as template** button. b. Enter a name in the *SAVE AS TEMPLATE* dialog. • From the 'Requests' tab or 'Stats for request page': a. Click the request's **save as template** icon. The **Create new template for** page displays.

Starting Tab or Page	Instructions
	b. Enter a name in the **Template name** field; then make any other changes following instructions in Creating a New Template. 2. Click the **Save** button. The template is added to the **Templates** tab.

Other Template Tasks

Viewing a Template Summary

Starting Tab or Page	Instructions
Templates Tab	1. Click the row containing the template or click the template's **preview** icon. The Template Summary page displays.

Editing a Template

Starting Tab or Page	Instructions
Templates Tab Template Summary Page	1. Click the **edit** icon. The *Edit template* page displays. All details copied from the template. 2. Make desired changes following instructions in Creating a New Template. 3. Click the **Save** button.

Duplicating a Template

Starting Tab or Page	Instructions
Templates Tab Template Summary Page	1. Click the **duplicate** icon. The *Clone template for* page displays with all details except the template name cloned from the template. 2. Make desired changes following instructions in Creating a New Template. 3. Click the **Save** button.

Deleting a Template

Starting Tab or Page	Instructions
Templates Tab Template Summary Page	1. Click the **delete** icon. 2. Click **Continue** in the dialog that prompts you to delete the template. The template is removed from the **Templates** tab.

Stats for request Page

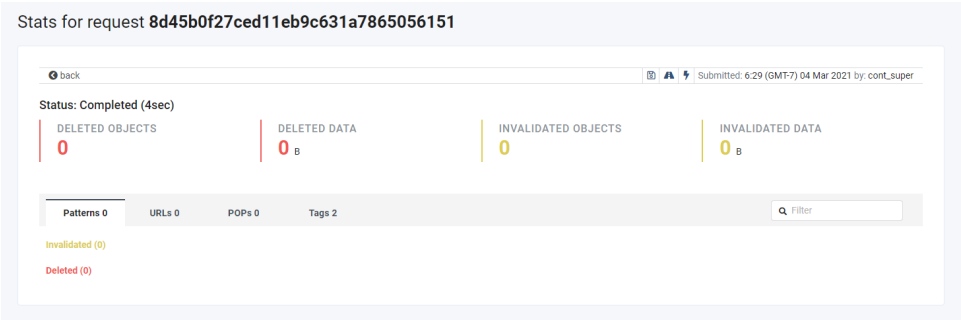

Stats for request pages display information about a dry run or purge.

Request ID

The request's id is displayed at the top left side of the page.

Toolbar

Text and icons on the right above the *Status* section provide information about the request and allow you to do additional tasks.

- template name: If the request was based on a template, the template name is displayed.
- **save as template** icon: save the request as a template. Available only in the **Stats for request** page.
- **dry run** icon: do a dry run based on the request.
- **run** icon: do a purge using the request.
- request information: date and time submitted, and submitter.

Status Section

Shows the request run time, number of deleted and invalidated objects, number of bytes for deleted, and invalidated objects.

Request Details Section

Selectable tabs allow you to view lists of items purged by patterns, URLs, and Tags. For longer lists, use the **Filter** field to locate the information you need.

Template Summary Page

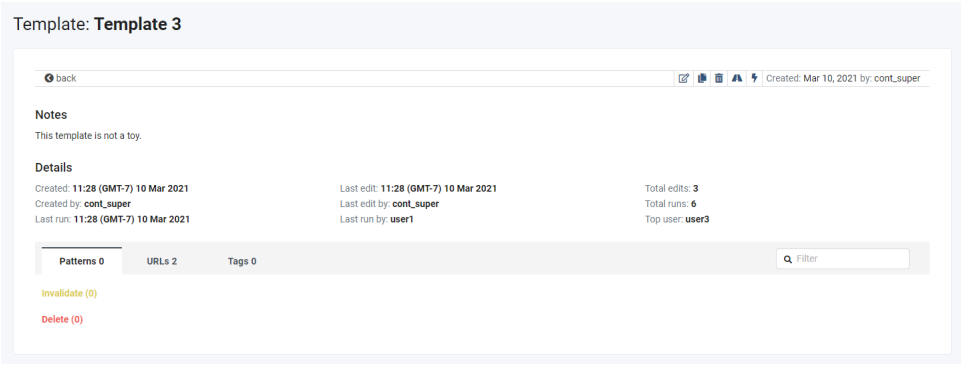

The Template Summary page provides read-only information about a template, including identifying information, notes, edits applied to it, and history of its use.

Icons at the top of the page allow you to edit or duplicate the template, delete it, do a dry run, or do a purge.

The *Details* section provides historical information about template creation and modification as well as run information (dry runs are not included). The section also shows items purged by patterns, URLs, and Tags. For longer lists, use the **Filter** field to locate the information you need.

> **Note:** For dry runs, the *Last run* and *Last run by fields* contain only a dash.

Purge Notifications

Purge notifications are displayed in the *Activity and Events* section of your account's dashboard.

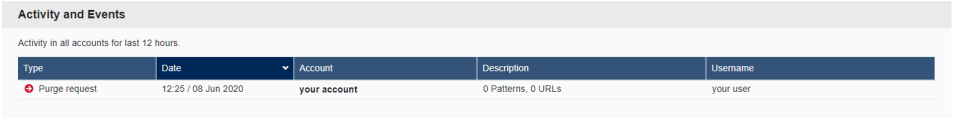

The SmartPurge REST API

The *SmartPurge* REST API provides programmatic access to all *SmartPurge* features. Some API features, such as unlimited callbacks, are available only with *SmartPurge Plus*.

For more information, please see the SmartPurge REST API User Guide.

SmartPurge Best Practices

- Consider the impact of deletions. Objects will be removed and cause cache fill for every purged object on the next request from a user. A large 'delete' purge can send a high amount of unwanted traffic back to your origin.
- When in doubt, invalidate. Invalidation hides objects from user access and causes an If-Modified-Since request HTTP header request against your origin, which will only cause cache fill only for objects that have changed.
- Don't submit more patterns than necessary. A wildcard at the root directory will reliably clear it out, so submitting secondary patterns for the same path will result in failed requests and are not needed. The *Stats for request dry run page* will verify that the purge successfully hit its mark.
- Wildcards can be used against Origin URLs only. Published URLs require a pattern with an exact match.
- Build your patterns for maximum effect, so you don't lose valuable time editing and resubmitting failed requests. Every purge request is sent to every server in Limelight's infrastructure.
- Use wildcards (asterisks) wisely and sparingly, especially in root directories, or when a bulk purge is not the intended outcome. Unexpected cache fills, lost files, rate limits, or failed purge requests can result.
- Know your workflow, especially if you find yourself juggling purge requests frequently. If your developers use versions on their CSS or JavaScript includes, you could lower their time to live (TTL) in the cache to better meet their needs. You can employ configuration techniques to decrease your need to purge. Reach out to your Solutions Engineer when in doubt.

Managing Authentication

You can configure Lightweight Directory Access Protocol (LDAP) and Security Assertion Markup Language (SAML) authentication types. This capability is intended for customers who use an SSO provider such as Okta or Ping Identity and allows customers to integrate Limelight Control authentication into their overall SSO capability.

By default, each company has Basic authentication enabled. Basic authentication is the standard type where users log into Control using their user, password, and 2FA authentication token. You cannot delete, edit, or deactivate the Basic authentication configuration.

The information in this section assumes you understand LDAP and SAML.

Notes:

- Only users with the Company Admin role can create configurations.
- You can create at most one configuration of each type (LDAP and SAML).
- The authentication capability is available only for companies that have the Authentication product.

Authentication List Page

After you navigate to Manage > Authentication, the *Authentication List Page* is displayed, and existing configurations are listed.

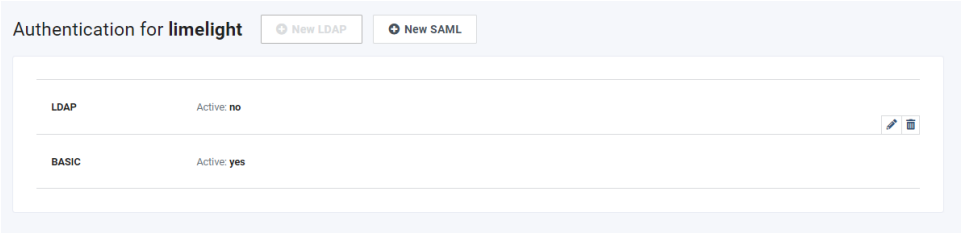

Creating a Configuration
Editing a Configuration
Testing a Configuration
Activating a Configuration
Deactivating a Configuration
Deleting a Configuration

Creating a Configuration

LDAP
SAML

Creating an LDAP Configuration

1. On the *Authentication List Page* click the **New LDAP** button.
   The *New LDAP* page is displayed.
2. Enter information in the fields (see Fields on the New LDAP Page).

3. Click the **Save** button to persist the configuration.

Fields on the New LDAP Page

Field	Description
Active	Indicates whether the configuration is active.
Base DN	Base Distinguished Name. Uniquely identifies the entry and its position in the directory information tree (DIT) hierarchy. Consists of the Organization Unit (ou) and Domain Components (dc).
Server URL	Address of the LDAP server that responds to client requests for information such as login credentials.

Creating a SAML Configuration

Limelight's SAML configurations are accessed over SSL to implement a secure connection.

1. On the *Authentication List Page* click the **New SAML** button.
   The *New SAML* page is displayed.
2. Enter information in the fields (see Fields on the New SAML Page).
3. Click the **Save** button to persist the configuration.

Fields on the New SAML Page

Field	Description
Active	Indicates whether the configuration is active.
Certificate (base64)	Certificate text, including the **-----BEGIN CERTIFICATE-----** and **-----END CERTIFICATE-----** markers.
Certificate fingerprint	The DER-encoded certificate's hash.
Fingerprint algorithm	Hashing algorithm that generated the fingerprint.
Idp Entity id	Globally unique name for the SAML entity, either an Identity Provider (IdP) or a Service Provider (SP).
Idp SSO url	Web address of the SAML IdP that handles sign-in requests.
User identifier attribute name	Defines the attribute to be sent by your SSO system to Control and used by Control to log in. Attributes vary depending on the SSO system, so this field allows you to define your own specific attribute name. As an example, assume the user identifier attribute name is 'EmailAddress'. On login, the SSO system sends an assertion request to Control with 'EmailAddress' = 'user-@mail.com'. Control will look for the 'EmailAddress'

Field	Description
	attribute in the assertion and try to authenticate the user with user@mail.com, which is the same as logging in to Control with login = user@mail.com on Control's login screen. Defaults to 'UserID'.
Private key	The unique string specific to you that you created when you requested the certificate with a Certificate Signing Request (CSR).
SSO URL (Assertion URL)	Required by some SSO Identity Providers when a configuration is being created, before the LDP SSO url has been issued. Users can copy and paste as needed into other fields such as **Idp SSO Id**.
SP Entity ID	Required by some SSO Identity Providers when a configuration is being created, but before the SP Entity ID has been issued. Users can copy and paste as needed into other fields such as **Idp Entity Id**.

Editing a Configuration

1. On the *Authentication List Page* click the configuration's edit (pencil) icon.
   The *Edit configuration* page specific to the authentication type (SAML or LDAP) is displayed.
2. Modify the fields on the page (see Fields on the New SAML Page and Fields on the New LDAP Page).
3. Click the **Save** button to persist the changes.

> **Note:** You can only edit LDAP and SAML configurations.

Testing a Configuration

1. Click the configuration's edit (pencil) icon on the *Authentication List Page*.
   The *Edit configuration* page specific to the authentication type (SAML or LDAP) is displayed.
2. Click the **Test Configuration** button at the bottom of the page.
   The TEST CONFIGURATION dialog is displayed.
3. Enter information in the dialog. The information requested depends on the authentication type.

Authentication Type	Fields
SAML	Username - a valid Control user name
LDAP	Control user email and password.

4. Click the **Test Configuration** button in the dialog. Test results depend on the authentication type.

Authentication Type	Results
SAML	Results open in a new browser tab.

LDAP	Results are displayed as a JSON object at the bottom of the dialog. Following is a sample success response:
	<pre>{
	"telnetSuccessful": true,
	"authSuccessful": true,
	"telnetTime": 175,
	"exceptionMessage": null
	}</pre>

Note: You can only test SAML and LDAP configurations.

Activating a Configuration

1. Hover the mouse pointer over a row on the *Authentication List Page*.

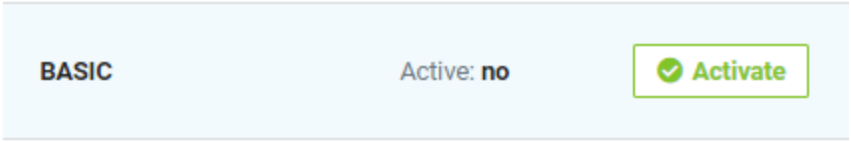

2. Click the **Activate** button.
   The Active status changes to **yes** and the **Activate** button label changes to **Deactivate**.

Note:
You can also active the configuration using the following steps:
1. Click the edit (pencil) icon for the configuration on the *Authentication List Page*.
2. On the *Edit configuration* page, put a checkmark in the **Active** checkbox.
3. Click the **Save** button.

Deactivating a Configuration

1. Hover the mouse pointer over a row on the *Authentication List Page*.

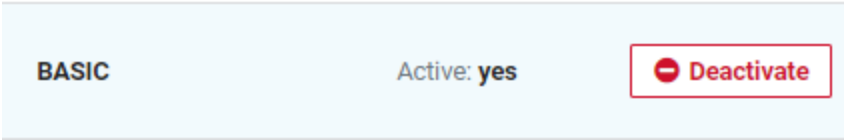

2. Click the **Deactivate** button.
   The Active status changes to **no** and the **Deactivate** button label changes to **Activate**.

Notes:
You cannot deactivate a configuration if it is the only active configuration.

You can also deactivate the configuration using the following steps:

1. Click the edit (pencil) icon for the configuration on the *Authentication List Page*.
2. On the *Edit configuration* page. remove the checkmark from the **Active** checkbox.
3. Click the **Save** button.

Deleting a Configuration

1. Click the configuration's remove (trash can) icon on the *Authentication List Page*.
2. Click the **Delete** button in the dialog that prompts you to remove the entry.
   Control deletes the configuration.

Notes:

- You cannot delete a configuration if it is the only active configuration.
- You cannot undo a deletion.

Managing Origin Storage Users

Navigate to Manage > Origin Storage Users in the navigation pane. The *Origin Storage Users* screen is displayed, with a list of all Origin Storage users associated with the currently selected Limelight Account.

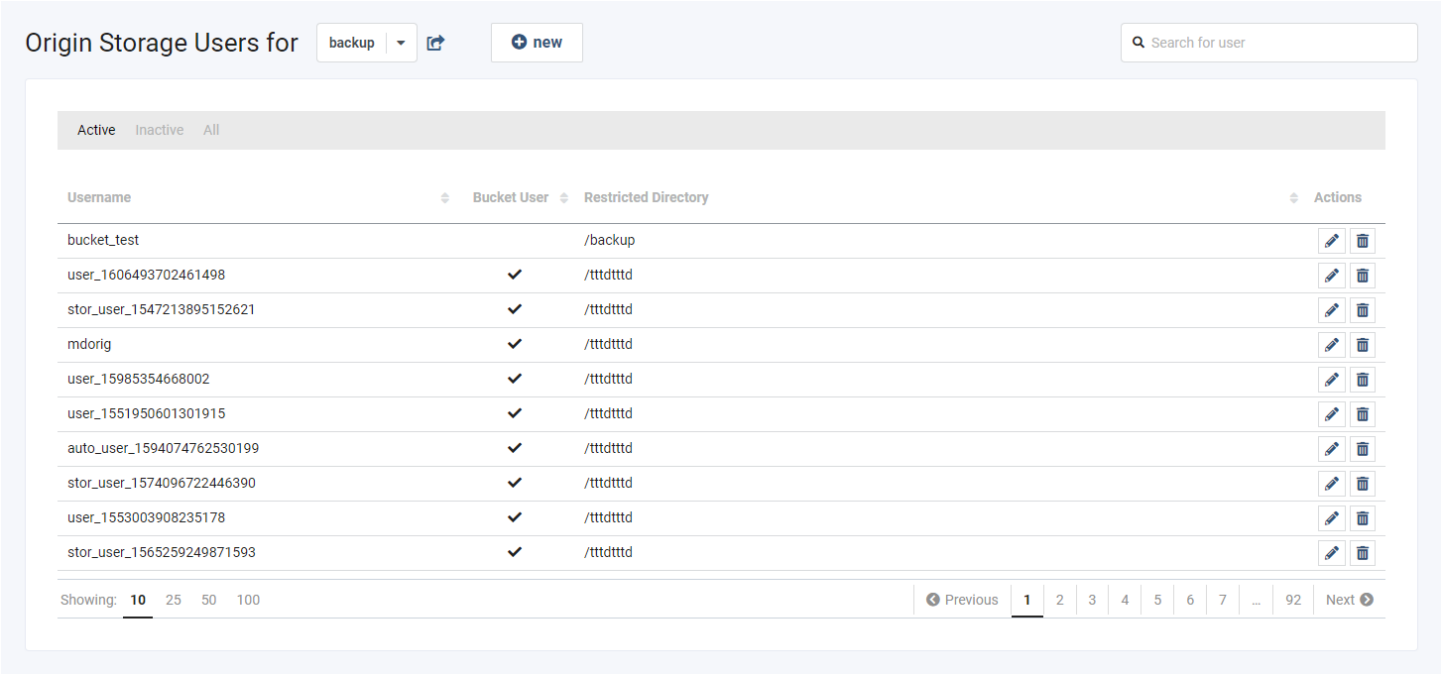

Links at the top of the list allow you to show active, inactive, or all users.

The list includes the following information for each storage account:

- **Bucket User** - Indicates if the user is allowed to work with buckets.
- **Username** - Displays the username of the user who created the storage account.
- **Restricted Directory** - Displays the path to the restricted directory.

Creating a New Origin Storage User Account

To create a new user Origin Storage account, click **+ new**, and the *Create Origin Storage User* page is displayed.

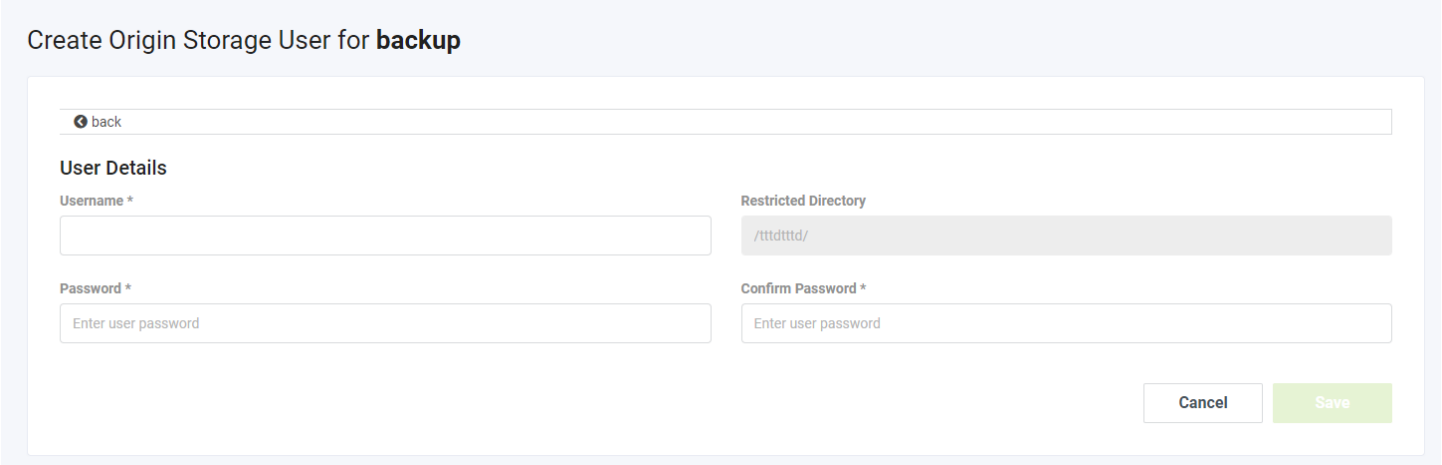

- Enter the **Username** of the user who will be managing this account
  Note: Once your Limelight Account has been saved, this field cannot be edited
- Enter the **Restricted Directory** path to your restricted content. For example: `/testtest/test`
  Note: The restricted directory is prepended with your shortname before the slash (/)
- Enter the **Password** and **Confirm Password** for this account
- Click **Save** to create the new account

Note:
If you are a bucket-enabled customer, you can only create bucket-enabled users, and the restricted directory is your bucket "root." It is pre-configured as `/<account name>_buckets` and is disabled:

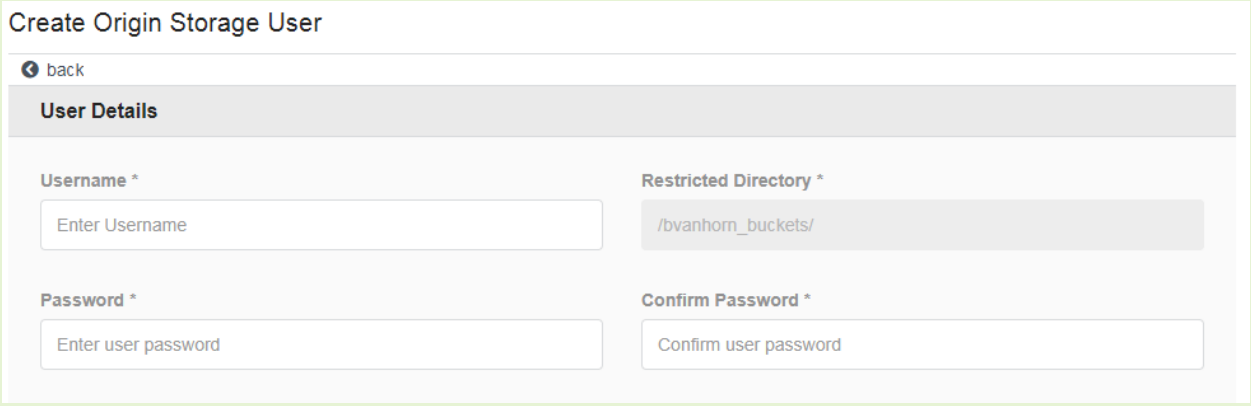

Otherwise, all UI elements and functionality for creating new user accounts are the same as for non-bucket-enabled customers.

Editing an Existing Account

In the *Origin Storage Users* screen, locate an existing account and click the pencil icon. The *Edit Origin Storage User* screen is displayed. The fields and controls for editing an account as the same as under Creating a New Origin Storage User Account, except that the **Username** and **Restricted Directory** fields cannot be edited.

Deleting an Account

In the *Origin Storage Users* screen, locate an existing account and click the trash icon, then click **Continue** in the confirmation dialog. The account will then be deleted.

Reactivating an Inactive User

Inactive users are shown in the **Inactive** and **All** tabs. To reactivate a user, click the **re-activate** (curved arrow) icon in the Actions column.

Exporting Origin Storage User Data

In the *Origin Storage Users* screen, click the **Export CSV** icon to the left of the **+ new** button.

User data is exported to a CSV file.

Managing Users

The *Users* page lets you add new users and manage existing users, providing extensive control over user permissions for portal features and report data.

Finding & Selecting Users

You can find specific users in three ways:

- Using the **Search for user** search field
- Filter users by last name by selecting one of the alphabet letters above the user list (**A**, **B**, **C**, and so on)
- Paging through the list of users with the numbered tabs, **Previous**, or **Next** buttons below the list

You can change the number of users shown in the list, click the corresponding number next to *Showing:* below the list.

To select a user, click the corresponding row in the user list.

Adding New Users

To add a new user, click the **+ new** button. The *Create User* page appears.

Fill out the fields on the page:

- **Username**
- **Email Address**
- **First Name**
- **Last Name**
- **Role**
- **Language**
- **New Password & Confirm Password**

To give the new user *Company Admin* privileges, use the **Role**drop-down menu. The default role is *User*.

Additional options:

- Use the **Two-factor authentication** checkbox to force the user to log into Limelight Control. See 2FA Security for additional information.
- To automatically generate a secure password for the new user, click the **Generate password** button.

When you are finished, click **Save**, and you will be returned to the *Users* page.

Migrating Users to Another Company

To migrate a user to another company:

1. Click the migrate icon on the right side of the user's row.
   The *MIGRATE USER* dialog is displayed.
2. Select the target company in the drop-down menu.
3. Click the **Migrate** Button.
   The user is migrated to the selected company and a success message is displayed.

Notes:

- Only users with the 'user' or 'company admin' role have the **Migrate** button available.
- Only users with migration permissions can migrate a user.

Editing User Profiles

To edit an existing user profile, including changing the user's password, click in the user's row or click the edit button at the right side of the user's row. The *User Details* tab is displayed.

> **Note:**
> Control does not enforce password expiration due to other security measures such as 2FA and transport over HTTPS; however, customers can change user passwords to align with their security practices.

In this view, an additional control, the **User Enabled** checkbox, becomes visible. If you need to revoke the user's access to the portal, disable (uncheck) the checkbox.

When you are finished, click **Save**, and you will be returned to the *Users* page.

Editing User Permissions

The *Permissions* tab allows you to manage user access to products, capabilities, and reports.

A list of the major groupings is displayed on the left side of the tab.

- *CONFIGURE*, *MANAGE*, and *REPORTS* refer to the corresponding menus in the navigation pane.
- *OTHER* contains permissions for more granular features.

Expandable items have a triangle to the left of the label, and items with a circle to the left can be individually configured.

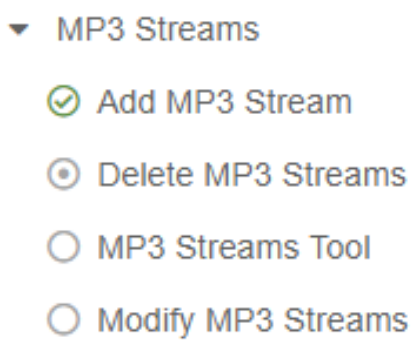

- A green checkmark in the circle indicates the item has been configured for all shortnames in the company[1] you selected after logging into Control.
- A dot within the circle indicates that the item has been configured for one or more or all shortnames but not the entire company.
- An empty circle indicates that the item has no configurations.

[1]The selected company is shown at the top right of the page.

When you select an individual item, the "path" (A) to the item is displayed on the right side of the screen along with (B) a grid of shortnames.

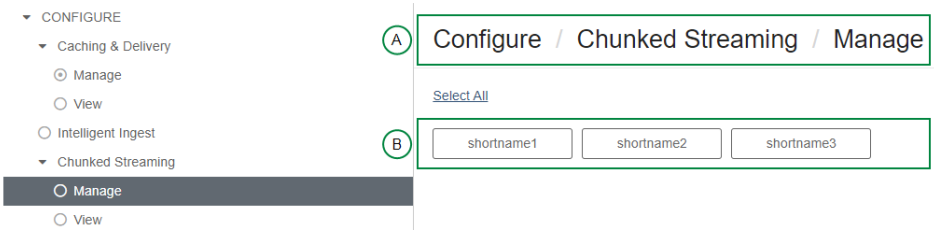

To enable configurable items:

1. Click an expandable item to drill down to configurable items.
2. Click a configurable item to reveal a list of shortnames on the right.
3. Configure the item:
   - Click the **All for** toggle to enable the item for the entire company.
   - Click the **Select All** link to enable the item for all shortnames. From there, you can deselect individual shortnames if desired.
   - Click individual shortnames to enable the item for only the shortnames selected.

Notes:

- Changes are saved immediately upon selection.
- If the **All for** toggle is not present, it means Limelight does not support per-company permissions.
- If later on a new shortname is added to the company, the shortname automatically appears in the shortname grids.
- By default, all users in a company have permissions to view the company's reports.

Cloning User Permissions

You can clone permissions from a "source user" to a "target user."

1. Locate the desired user (see Finding and Selecting Users).
2. Click the **clone permissions** icon

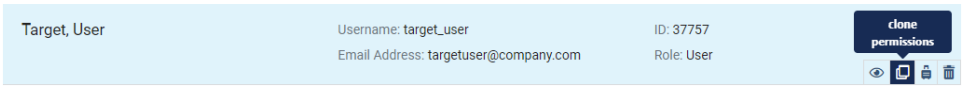

The **CLONE USER PERMISSIONS** cialog is displayed.

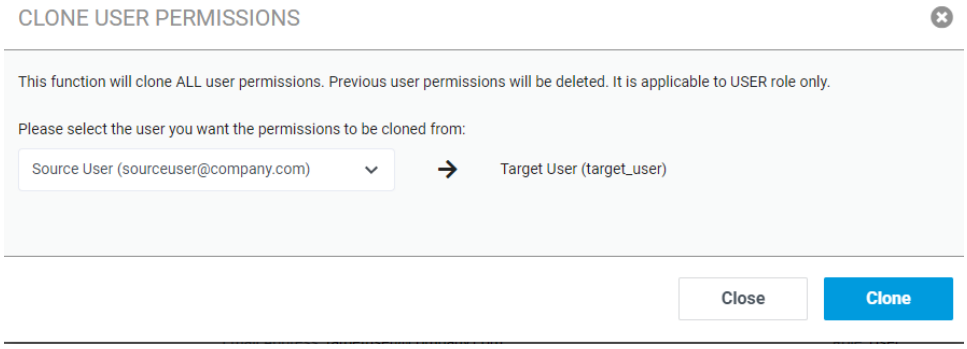

3. Select the source user in the drop-down menu; then click the **Clone** button.
   A success message is displayed.

Notes:
- Only company admin users can clone permissions.
- Both source and target users must belong to the same company and must both have the USER role.
- The target user's permissions are replaced with the new permissions.

2FA Security

As part of Limelight's commitment to maintain the highest level of security for our customers, we use *Two-Factor Authentication (2FA) with Limelight Control*.

This capability is for all *Control* users with either:

- the *Company Admin* role, or
- *Manage* permissions for *Configure*

When *2FA* is enabled, the *2FA* status of your users is displayed at the top of the *User Details* tab.

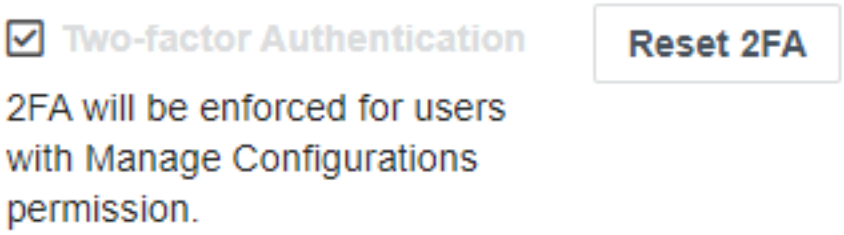

Users with the *Company Admin* role can enable or disable *2FA* for individual users by setting the **Two-factor Authentication** checkbox.

Control users who are enabled for *2FA* must enter a new *2FA* security code each time they log in to *Control*. Users can quickly generate *2FA* security codes using a mobile app (such as the Google Authenticator App). Once an authenticator app is set up, a network connection is not required to generate new tokens.

Free *2FA* authenticator apps can be found here:

- for iOS
- for Android
- for Windows phone
- for Blackberry

User Experience

When *2FA* is enabled for *Control* users, they will be presented with the *2FA Get Started* screen the next time they attempt to log in. The screen includes an explanation of *2FA*, and a list of the steps users must follow to obtain a security code:

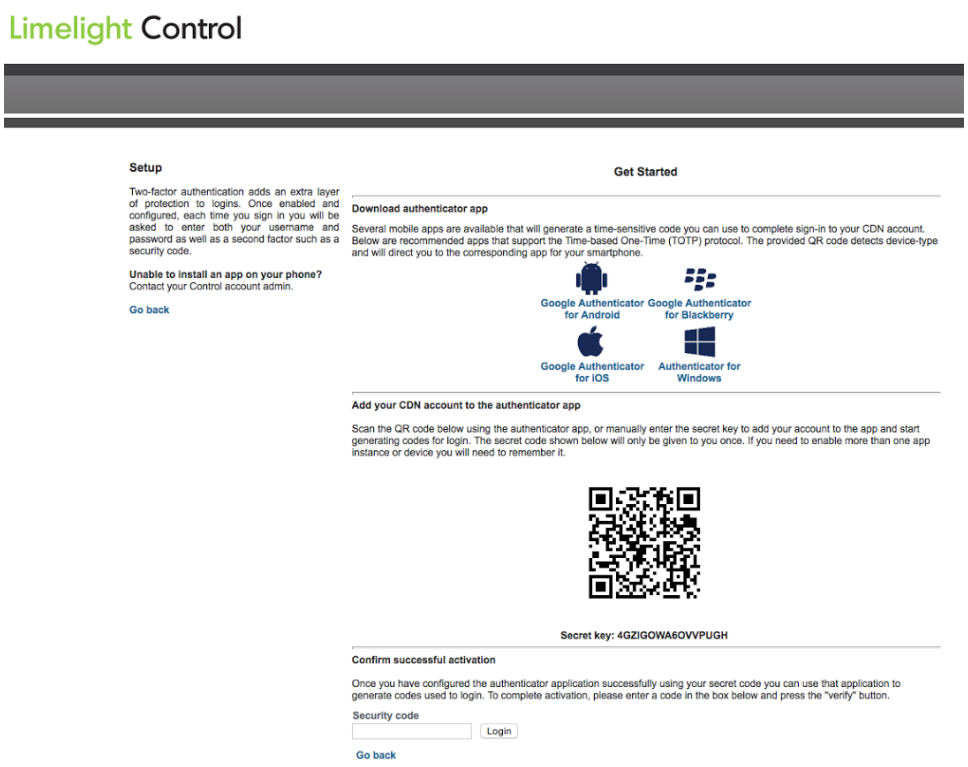

The steps are:

1. **Download App**. On a mobile device, download a supported authenticator app:

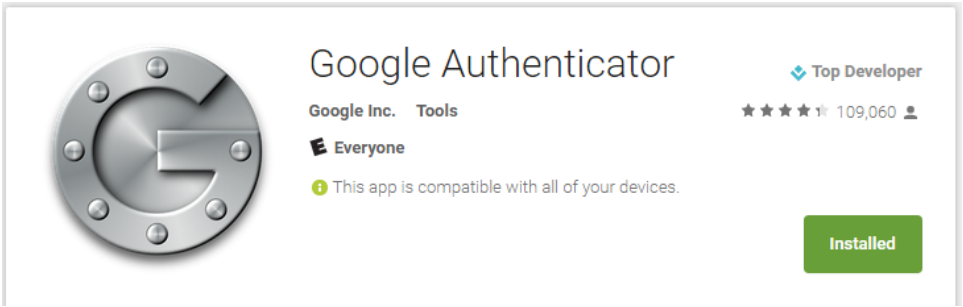

2. **Add Account**. In the authenticator app, follow the app-specific instructions to add the *Limelight Control*. For example, in the Google Authenticator App, the *Scan a barcode* option can be used to automatically add the portal by capturing the QR code in the *2FA Get Started* screen:

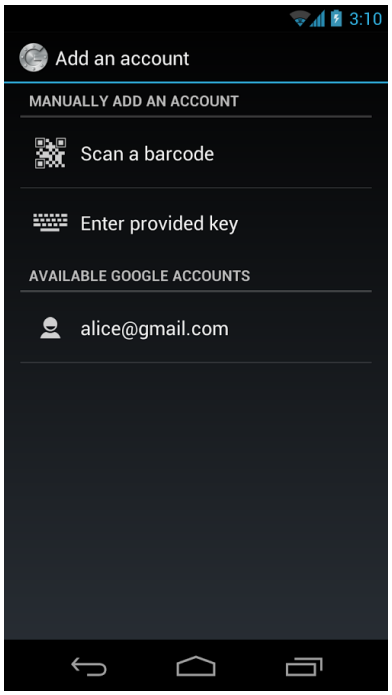

Or manually enter the *Secret Key* (under the QR code) using the *Enter provided key* option.

Note: When setting up *2FA*, *Limelight* recommends users record the *Secret Key* displayed by the *2FA Get Started* screen. The code allows the configuration of multiple devices and browser authenticator apps so that if a user's smart-phone is unavailable, they can use another paired device or app. The *Secret Key* should be kept safe, just like any other password.

3. **Obtain Security Code**. In the authenticator app, follow the app-specific instructions to get a new *2FA* security code for *Control*:

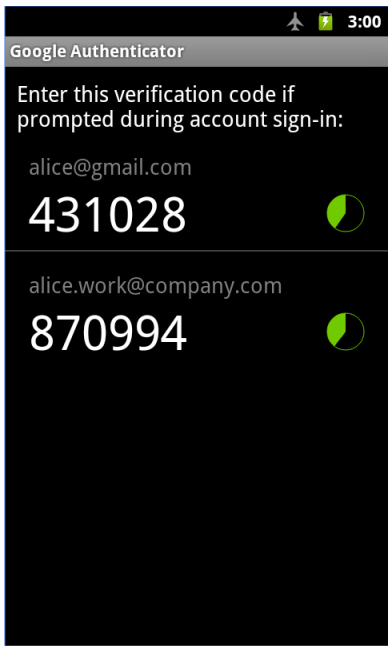

Once set up, the authenticator application will display a six-digit security code, which changes every 30 seconds.

4. **Enter Security Code**. In the *Control2FA Get Started* screen (or the *2FA Token* screen on subsequent use), enter the security code in the *Security Code* field and click *Verify*. If valid, the *Login* screen is displayed.
   - *2FA Get Started* screen:

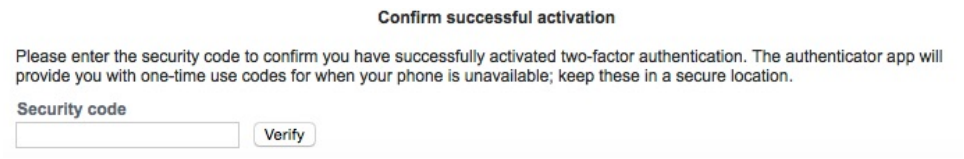

   - *2FA Security Code Confirmation* screen:

5. **Login**. In the *Login* screen, enter credentials and click *Login*

Resetting Device Pairings

Control users with the *Company Admin* role can reset device pairings for those with the *User* role using the **Reset 2FA** button:

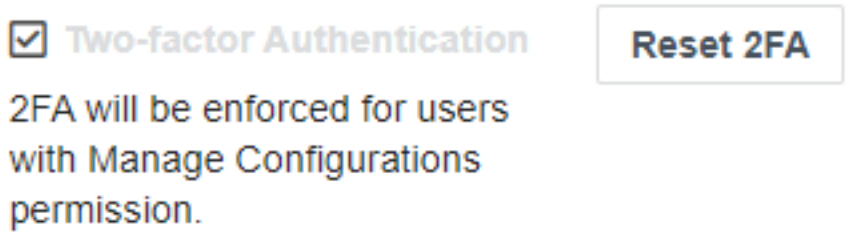

After a reset, the *Control* user must repeat steps *2* and *3* above to create a new pairing. This feature may be useful when users need immediate access, but their paired mobile device is not available. **Note**: This will reset the pairing for all currently paired devices for the user.

Company Admin users cannot use *Control* to reset device pairings for themselves or other *Company Admin* users. Resets for *Company Admin* users must be requested directly from Limelight Customer Service.

Other *Control* users should contact their *Company Admin* to request a reset.

Frequently-Asked Questions

What happens if a Control user upgrades their phone to a new model?

A reset is required. The user must download the authenticator app to the new smartphone and contact a *Company Admin*, who can then follow the steps in the Resetting Devices section. The user must then follow steps *2* and *3* in the User Experience section to pair with the new device.

What happens if a user forgets their phone or their battery dies?

Unless the user has recorded the *Secret Key* and is prepared to pair with an additional device or app (see next question), a reset is required. Only the currently-paired device and authenticator app can generate a code without performing a reset. At this time, we do not have an alternate means of providing codes. Users must have the original device and authenticator app or reset and re-activate with a new device.

Can a user 'pair' Control with multiple devices and authenticator apps?

Yes, if they record and then later use the *Secret Key* provided during the pairing process. In the *2FA Get Started* screen at the beginning of the pairing process, the *Secret Key* in a long alphanumeric string under the QR code. If recorded, the *Secret Key* can be used to pair with authenticator apps on multiple devices. Otherwise, it is not possible to pair with additional devices. **Note**: the *Secret Key* is sensitive data and should be guarded like a password.

What do users do if they don't have a "smart" mobile device?

Since the Google Authenticator algorithms are commonly known, several developers have also published authenticators for browser extensions and apps. These should work with *Control* as long as they use the same time-based OTP algorithm used by Google Authenticator.

If a user is granted Manage permissions for Configure, is 2FA also enabled?

Yes. Granting a user any of the *Manage* permissions for Configure automatically enables *2FA* for that user. This additional security is automatically applied because users with *Manage* permissions can add, change and delete product configurations, potentially affecting the delivery of production content.

Does 2FA affect access to Control APIs?

No, *2FA* only affects user logins via the *Limelight Control*. At this time, API access is secured using a separate shared secret and is not affected by *2FA*.

If you need to revoke a user's access to the portal, select the user, and in the *User Details* tab, uncheck the ***User Enabled*** checkbox. You can re-enable access at any time.

Origin Storage Console

Origin Storage Console Overview

The Origin Storage Console (the Console) is the user interface to *Origin Storage* that lets you perform many of the same tasks you might do through APIs, such as uploading files, creating directories and subdirectories, and so on. Depending on the speed of your internet connection and file size, you can populate your selected storage locations around the world in minutes.

If you already have an *Origin Storage* user account, you can access and use the Console; it mirrors your content. You can still access your *Origin Storage* account via the API.

If you are new to *Origin Storage*, you have the option to use either (or both) the Console or the *Origin Storage* API to store data in Limelight's many geographic locations around the world that best meet your business requirements.

> **Note:** The Console is intended for ad-hoc work only and should not be part of normal CDN work flows. Whenever possible, we recommend that customers leverage the Origin Storage API due to its feature-rich ingest workflow and its superior performance.

Console Workspace

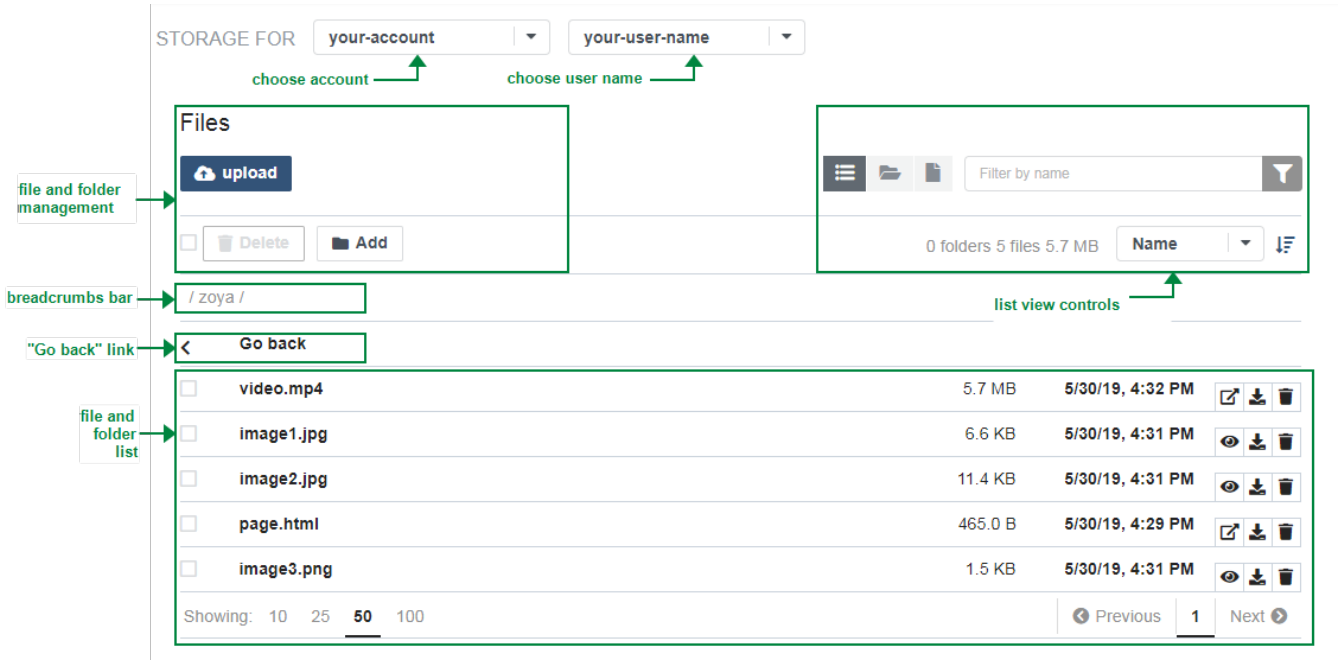

The Console Workspace contains these controls:

- **choose account / choose name**: Allow you to choose the account and user name
- **file and folder management**: Contains controls for file- and folder- management tasks such as uploading and deleting files and creating folders. (You can also delete files while viewing them in the files and folders list.)
- **list view controls**: Allows you to filter the content of the file and folder list.
- **breadcrumbs bar**: Shows your place in the content structure and allows you to navigate by clicking path segments. Each segment in the breadcrumbs bar is a link to that segment in the folder structure. Click a link to move to that location in your content structure.
- **Go back link**: Visible only when you are one or more levels from the root folder, the link takes you to the previous path segment listed in the breadcrumbs bar.
- **file and folder list**: Shows files / folders contained in the folder currently shown in the breadcrumbs bar.

Working with the Files and Folders List

Before You Start

Select an account and user name if you have not already done so. Use the drop-down menus at the top of the workspace:

Filtering

You can filter the contents of the files and folders list by typing text in the **Filter box** in the top right corner of the page.

> **Note:** The filter applies only to the folder you are currently viewing as reflected in the breadcrumbs bar, and not the entire directory structure of your content as a whole.

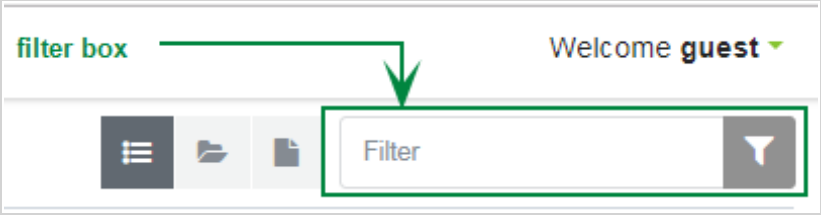

The Console matches files and folders as specifically as possible to your filter. The Console applies the filter to file and folder names.

As you type, the Console modifies the files and folders list to match your filter.

Filtering is case insensitive.

Controlling List Content

Using the following list options icons in the top right corner of the page, you can view:

- files and folders
- folders only
- files only

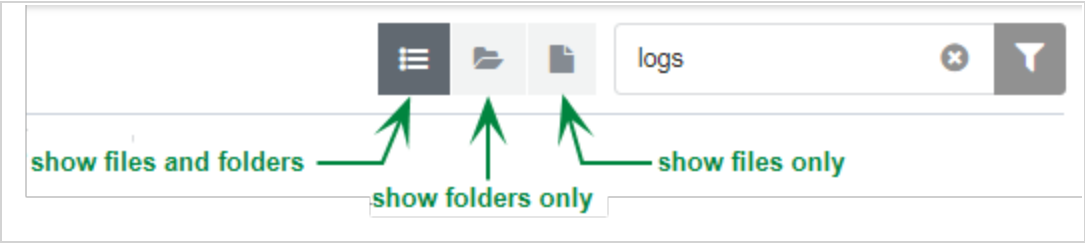

Sorting

Using the **sort drop-down** and **sort icon**, you can arrange the contents of the files and folders list.

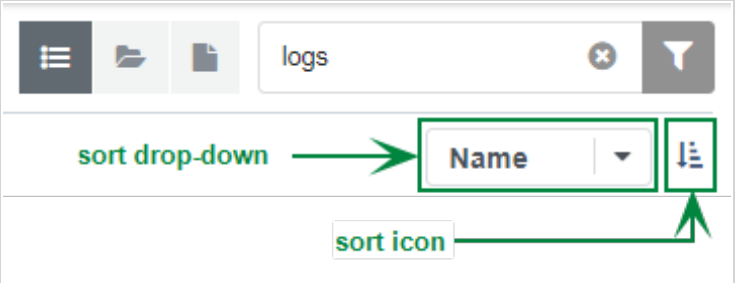

Choose a field to sort by in the **sort drop-down**:

- **Name**
- **Date Uploaded**
- **File Size**.

For folders, **Date Uploaded** is the folder creation date.

Click the **sort icon** to toggle between ascending (default) or descending order.

Paging

Use the **results per page** links at the bottom of the page to control the number of items visible per page. The **paging links** allow you to page through the **files and folders list** by selecting the page number you want to view.

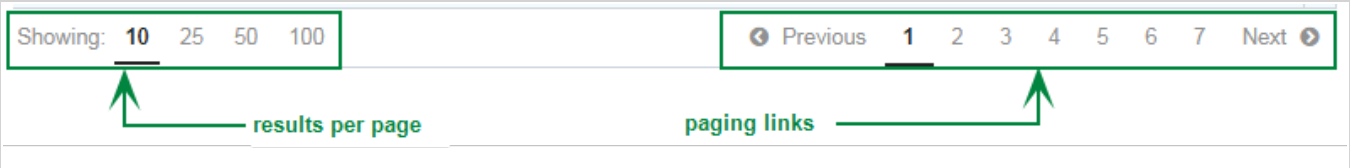

Working with Files

Before You Start

Select an account and user name if you have not already done so. Use the drop-down menus at the top of the workspace:

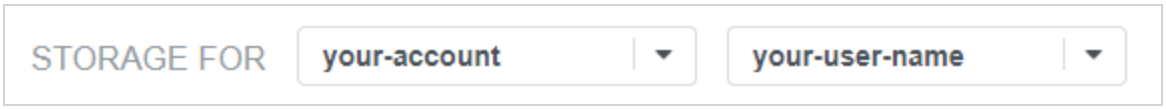

Deleting Files

You can delete a single file or multiple files.

Deleting a Single File

1. Navigate to the folder that contains the file you want to delete.
2. Click the **delete icon**.

3. Click **Delete** in the dialog that asks you to confirm the deletion.
4. The Console deletes the file.

> **Note:**
> You can also delete a file using the *Origin Storage* deleteFile API.

Deleting Multiple Files

1. Navigate to the folder that contains the files you want to delete.
2. Choose files you want to delete. You have two options:
   - Individual: Select individual files by clicking the checkboxes to the left of the file names.
   - All: To select all files currently visible in the **files and folders** list (based on the results per page setting), Click the checkbox to the left of the **Delete** button.

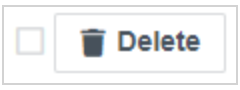

> **Note:** The button to delete files is inactive until you select one or more files.

3. Click the **Delete** button.
4. Click **Delete** in the dialog that asks you to confirm.

Previewing Images

You can view image files that are in your Origin Storage account.

1. Navigate to the folder that contains the file.
2. Click the **preview icon**.

3. The image appears in a dialog along with its direct link.

Getting Direct Links to Files

The Console allows you to obtain a file's URL that you can share by copying and pasting. The URL allows users to directly download the file from *Origin Storage*. For security purposes, the link is valid for 24 hours only.

1. Navigate to the folder that contains the file.
2. Click the **generate link icon**.

3. The direct link appears in a dialog box along with the link expiration date/time.

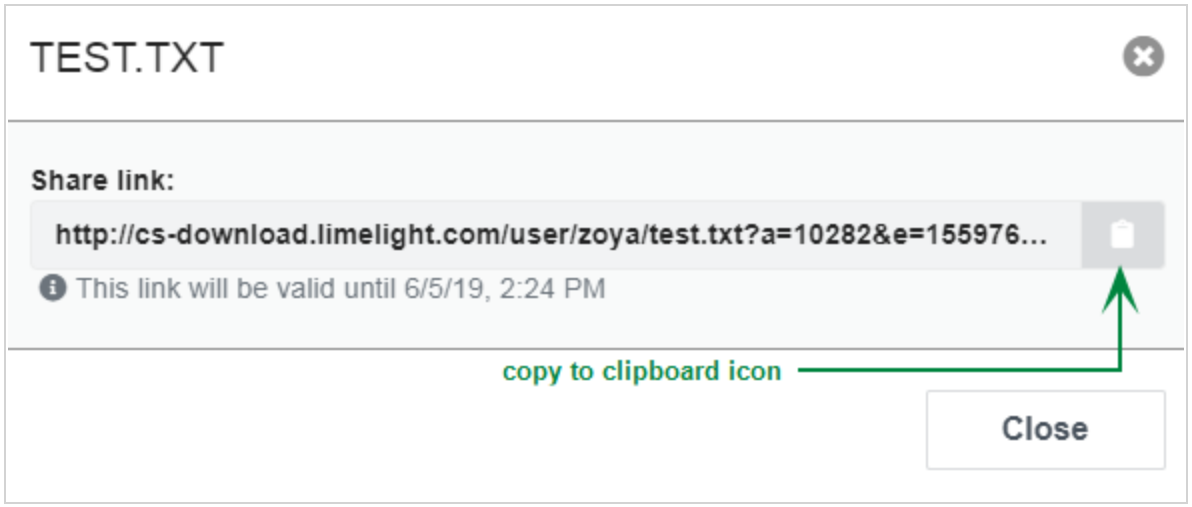

4. Click the **copy to clipboard icon** to get a link you can paste.

> **Note:**
> The *Origin Storage* API does not provide a means for obtaining a file's URL.

Downloading Files

1. Navigate to the folder that contains the file you want to download.
2. Click the **download icon**.

Your browser downloads the file.

Uploading Files

You can upload by drag and drop or by selecting individual files. The Console fully supports file names with UTF-8 characters.

> **Notes:**
> The Console allows you to upload a zero-byte file and imposes no file size limits.
>
> The Console allows you to upload no more than 50 files at once. At any given time, the Console uploads only two files concurrently. The rest of the files are in an upload queue.
>
> If you attempt to upload a file in a directory that already contains the file name, the Console over-writes the existing file.
>
> There is no limit to the number of files that you can store in a directory, but the Console displays a maximum of 10,000, so anything over that limit will not visible within the Console.

Uploading by Drag and Drop

1. Navigate to the folder to which you want to upload files.
2. Select one or more files from your desktop and drag them over the Console. As you hover over the Console, you see the *Drop files here* prompt:

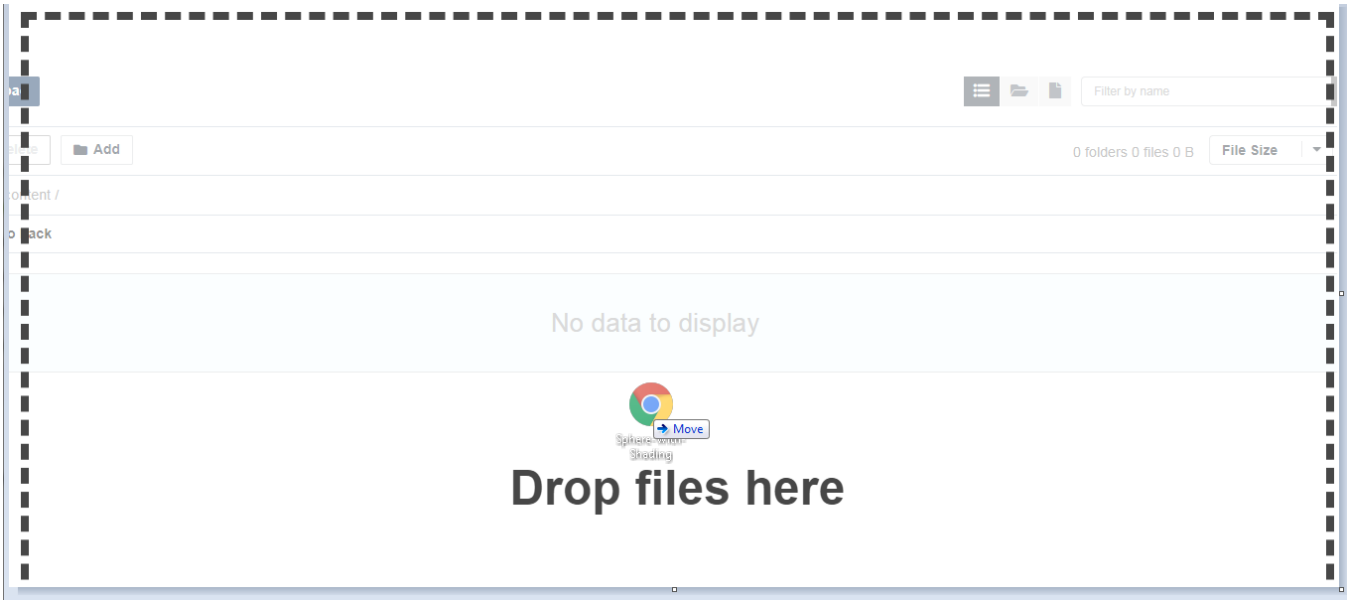

3. Release the files.
4. The Console uploads your files.

See also Viewing Upload Progress and Canceling Uploads.

Uploading Using the Upload Button

1. Navigate to the folder to which you want to upload files.
2. Click the **upload** button.

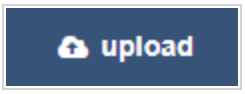

3. The **Open** dialog appears. Use the dialog to browse to the location of your files and select the desired files.
4. Click the **Open** button in the dialog.
5. The Console uploads your files.

See also Viewing Upload Progress and Canceling Uploads.

Viewing Upload Progress

During file upload, the Console displays the **upload progress** bar and the **Cancel All** icon:

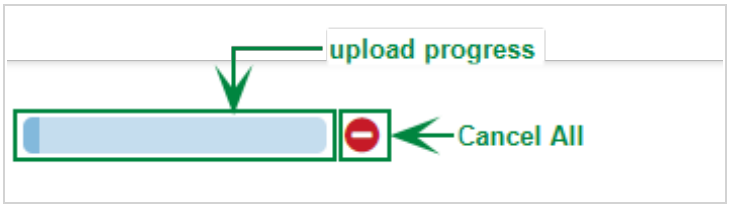

> **Note:** Small files are generally uploaded very quickly and the progress bar goes away almost immediately.

If you are uploading multiple files, you can click the **upload progress** bar to show *upload details*–the status of each file in the upload:

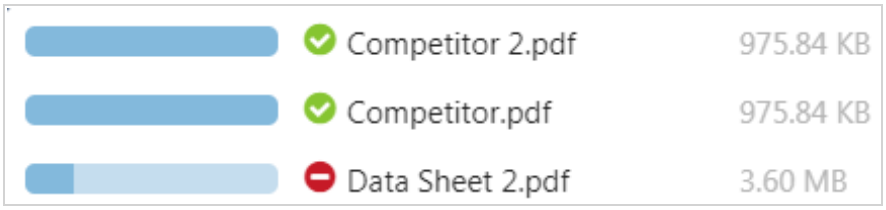

For uploads with more than five files, the *upload details* includes the **Show all uploading** link:

Click the link to view details for the remaining files.

Identifying Newly Uploaded Files

After the Console uploads a file, the Console flags the file with the *New icon*:

The file is placed at the start of the files and folders list.

If you don't see your uploaded files in the files and folders list, the list might be set to **show folders only**. Click the **show files and folders** button or the **show files only** button (see Controlling List Content).

Canceling Uploads

You can cancel a single file upload or cancel all files in a multi-file upload by clicking the **Cancel All** icon:

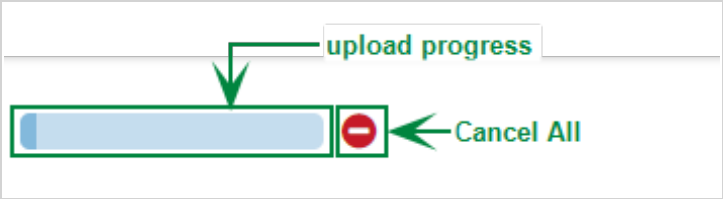

You can also cancel individual files in a multi-file upload:

1. Click the **upload progress** bar to show progress details.

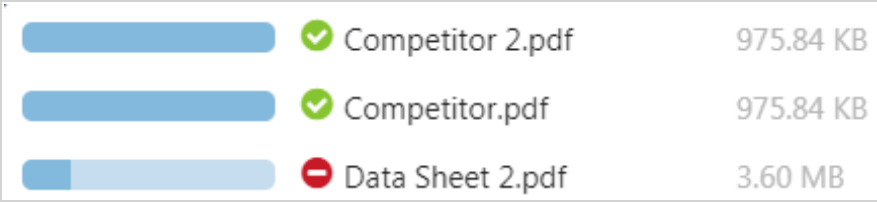

2. Click the **Cancel** icon for any individual uploads you want to stop.

Working with Folders

Before You Start

Select an account and user name if you have not already done so. Use the drop-down menus at the top of the workspace:

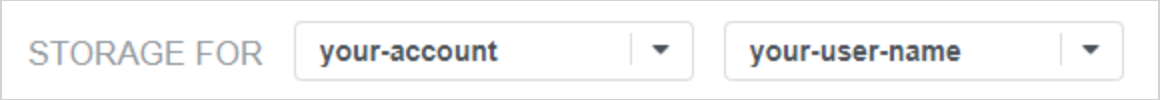

Creating Folders

1. Navigate to the folder where you would like to create the new folder.
2. Click the **Add** button.

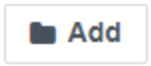

3. Enter a name in the *CREATE* dialog and click **Create**.

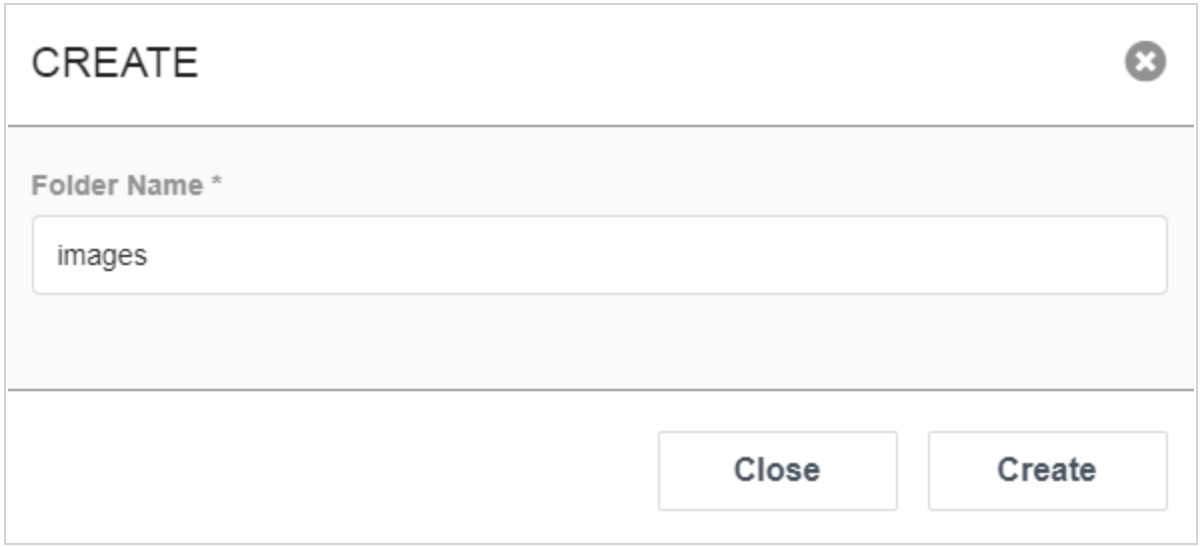

Note:
If you attempt to create a directory with a / slash in the name one of two things happens:

1) If the folder name before the slash exists, the Console creates a new directory in that folder.

2) If the folder name before the slash does not exist, the Console displays an error informing you that the parent path does not exist.

For example, if you are in the root directory and a folder named `test` exists and from the root directory you click **Create Folder** and enter `test/sub-test`, the Console creates the folder `sub-test` under `/test`.

Note:
You can also create directories using the *Origin Storage* API. Use any of the following:

- makeDir
- makeDir2
- post/directory

Uploading Folders

You can upload one or more folders by drag and drop.

- All files within the folder will be uploaded.
- You can upload multiple folders at the same time.
- You can upload a folder that has multiple levels of subfolders. All subfolders and their contents will be uploaded as long as the total number of files is less than a configured maximum number.

To upload folders:

1. Navigate to the folder to which you want to upload folder.
2. Select one or more folders from your desktop and drag them over the Console. As you hover over the Console, you see the *Drop content here* prompt:

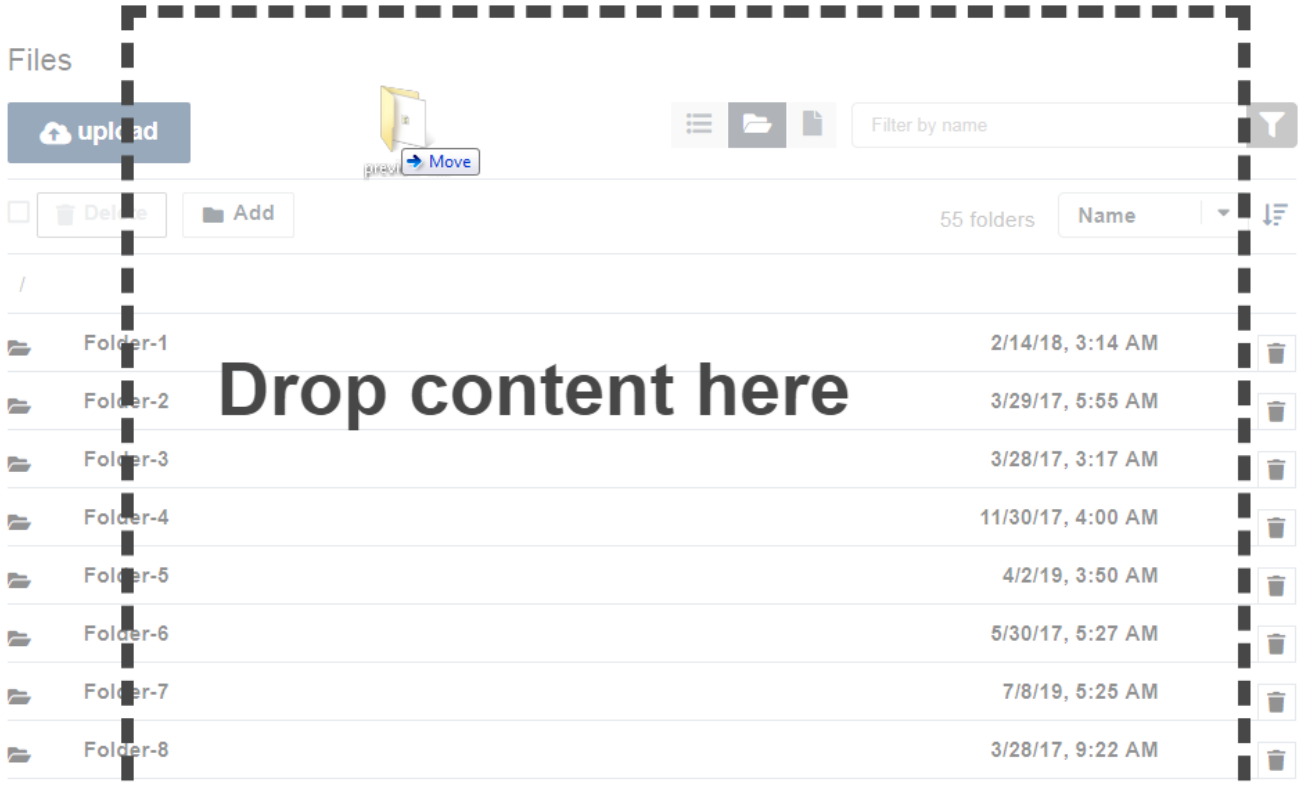

3. Release the folders.
4. The Console uploads the folders.

Note:
Visibility of uploaded folders and files depends on the view in which you did the upload:

- 'folders only' - Uploaded folder(s) are visible but not the files within. You need to change to one of the other filter views.
- 'files only' - Uploaded folder(s) are not visible. You need to change to one of the other filter views.
- 'files and folders' - All uploaded folder(s) and the content within are visible if navigated to.

Deleting Folders

Notes:

- You can only delete empty folders. To delete files in a folder, see Deleting Files.
- The Console does not support recursive deletes (deleting a folder's subfolders, all their subfolders, and so on).
- You can also delete a folder using the *Origin Storage* deleteDir API.

1. Navigate to the folder that you want to delete.
2. Click the **delete icon**.

3. Click **Delete** in the dialog that asks you to confirm the deletion.
4. The Console deletes the folder.

Managing Accounts and Users

> **Note:** Except where noted, the procedures in this section require Company Admin privileges.

Creating Origin Storage Console Users

1. Navigate to Manage > Origin Storage Users in the navigation pane.
2. Click **+new**, and enter the appropriate Origin Storage credentials and restricted directory.

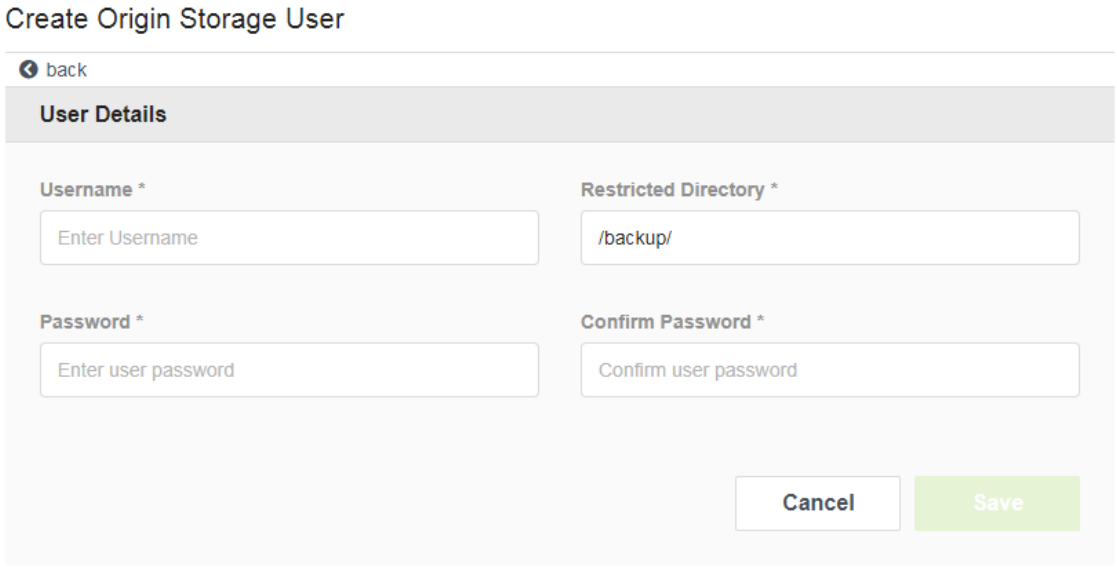

3. Click **Save**.

Granting Origin Storage Console Access to Existing Control Users

> **Note:** This procedure requires Company Admin privileges. If you have non-admin privileges, use the procedure in Associating Storage Users with the Origin Storage Console.

1. Navigate to Manage > Users in the navigation pane.
2. Locate the user in the directory.
3. Click the **eye** (user details) icon on the right side of the user's row.

4. Select the **Storage Users** tab.

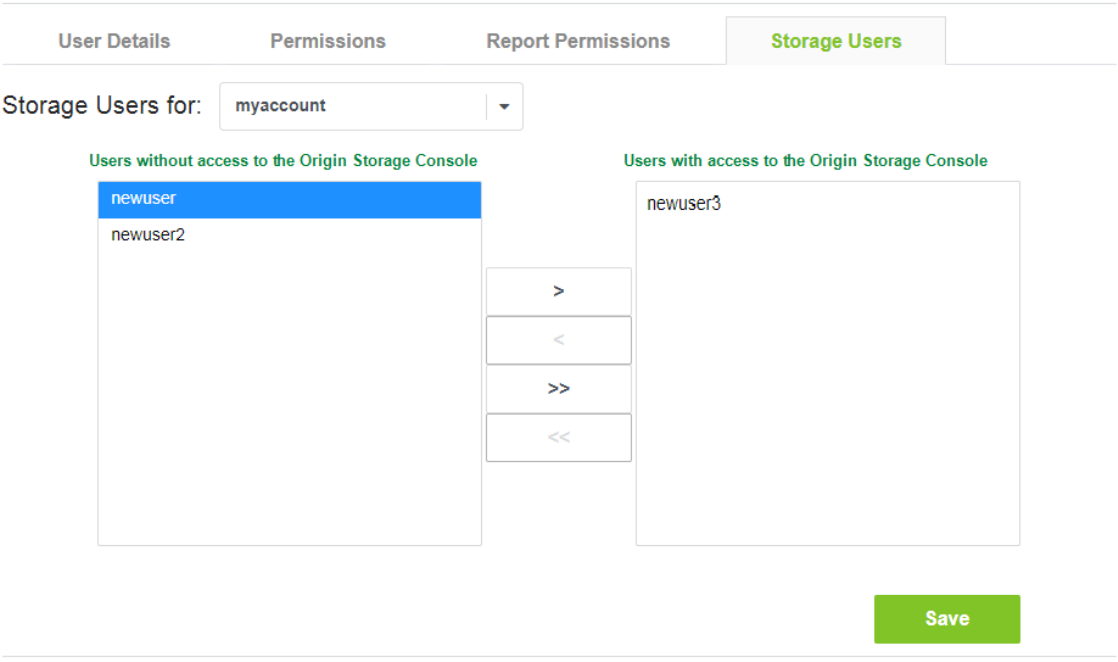

5. Use the controls on the page to grant access to the desired users.
6. Click **Save**.

Associating Storage Users with the Origin Storage Console

> **Note:** This capability is designed for non-admin users with the goal of enhancing the process of migrating external Origin Storage Console users to the Storage Console within Control.

If you want to allow a non-Control Portal user to access the Storage Console within Control and not access any other Control capabilities, you can associate the user with the Storage Console.

To associate a user:

1. Navigate to Manage > Origin Storage Console in the navigation pane.
2. Click the **Associate Storage User** button at the upper left corner of the screen.
3. In the subsequent dialog, enter the user's user name and password .
4. Click the **Add** button.

The user can now log into Control and access only the Origin Storage Console.

Changing Origin Storage Console User Passwords

1. Navigate to Manage > Origin Storage Users in the navigation pane.
2. Select a user.
3. Enter and confirm the new password.

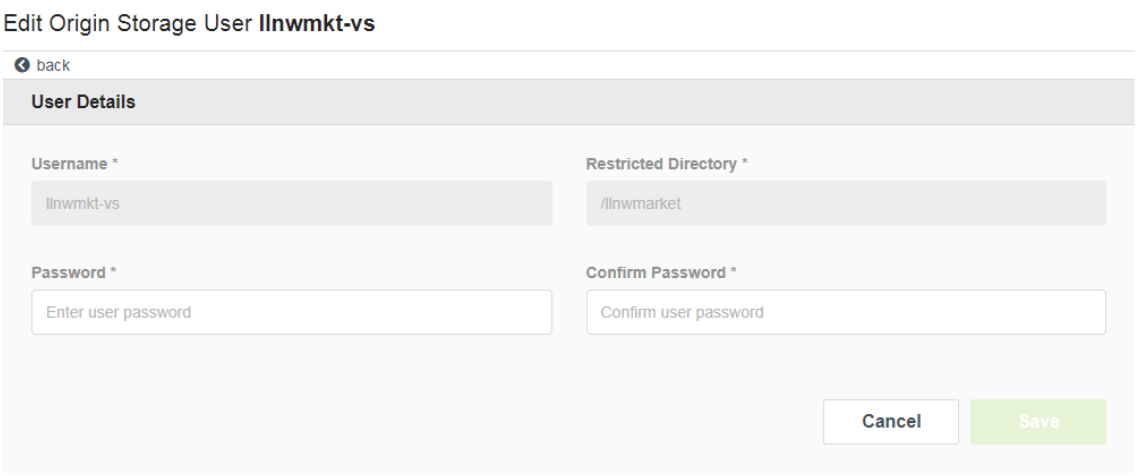

4. Click **Save**.

Logging Out

Click the **Welcome** link at the top right part of the Console and select **Logout**:

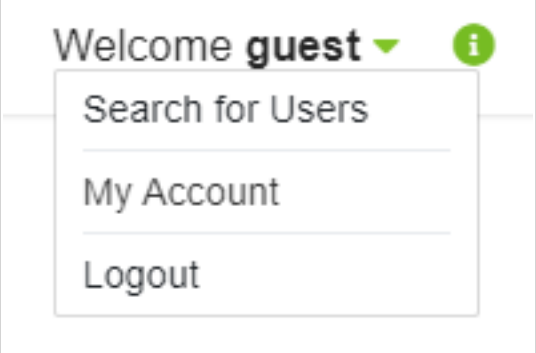

Viewing IP Allow Lists

You can view the current IP addresses of LimelightEdge Servers to update your firewall, without logging into the Limelight Control. The IP allow list returned from the API call is versioned, so you can compare the current version with the prior version to see how the list has changed.

You can use either the REST API or the RSS feed to view the IP allow list.

API Specification

The endpoint returns a versioned list of allowed IP addresses, expressed in JSON.

HTTP Method: GET

URL: `https://control.llnw.com/aportal/api/ipam/getIpAllowList.do`

Required Request Headers

Header	Description
`X-LLNW-Security-Principal`	Caller's user name
`X-LLNW-Security-Timestamp`	Current system UTC time in milliseconds
`X-LLNW-Security-Token`	HMAC-256 digest calculated using the caller's API shared key on: HTTP Method + URL + timestamp

Sample Request

`https://control.llnw.com/aportal/api/ipam/getIpAllowList.do`

Sample Request Headers

`X-LLNW-Security-Principal: sample_user`

`X-LLNW-Security-Timestamp: 1465226821474`

`X-LLNW-Security-Token:b62f93cdde95e94b814ba824430a25cfd31fc13f485f201f6878754caf6f0493`

Sample JSON Response

```
{
       "ipAllowList":["1.9.58.160/27",
       "37.238.255.224/27",
       "41.63.64.0/18"],
       "version":1
}
```

RSS Feed

Subscribe to the following:

`https://control.llnw.com/aportal/support/documentation/iprssfeed`

Each IP address is in a `content:encoded` element within a parent `item` element. Example:

```
<item>
       <content:encoded>69.28.128.0/18</content:encoded>
</item>
```

You can also view the recent IP allow list in customer documentation. Navigate to Help Center > Documentation in the navigation pane. Then select IP Allow list in the green navigation panel.

Reporting

Introduction

The *Reports* section of the *Limelight Control* gives you insight into the performance of each of your Limelight services. Summary and detail views are available for a variety of performance and billing-related metrics.

> **Note:** Some reports must be purchased or specifically enabled by Limelight before they become available in the portal. Please contact your Account Manager or Limelight Support if you have questions about access to a specific report.

Reports General Information

Sources of Inconsistency Between Control Reports

Three independent systems collect data for Control Reports:

System	Description
EQ Billing	Processes data for the Billing report.
EQ Reporting	Processes data for these reports: • Status Codes • Live Stats • Traffic (all tabs except **Hosts & URLs**). The EQ Reporting system is optimized for low-latency data processing.
RLDP	Generates data for the **Hosts & URLs** tab on the Traffic Report. The RLDP system uses raw log data. Data for the **Hosts & URLs** tab is refreshed once a day at about 0900 am GMT-7.

Due to processing differences, the resulting metrics differ across the three systems. Various network and hardware issues can also affect data collection and processing that occurs in Limelight's distributed environment. As a result, customers might see a small metrics deviation, usually less than 1%.

Request Proration

Proration is enabled only on the Billing Report and Traffic Report (**Hosts & URLs** tab only). Proration makes data aggregation significantly (a few times) more expensive.

Often, customers' reports show a fractional number of requests because requests are not always instantaneous. Due to the large media downloads and streams that Limelight serves on its network, requests normally take a significant amount of time to complete.

To account for this, Limelight spreads out requests throughout their lifetime. For example, a customer has a 500 megabyte file that users download. A user starts downloading the file at 8pm on Tuesday, and the download finishes at 1am on Wednesday. This is the only traffic that the customer has during that time.

Since 4/5ths of the file was downloaded on Tuesday and 1/5th of the file was downloaded on Wednesday, the reports will show the following data during that time:

Day	Requests	Bytes	Seconds
Tuesday	0.8 (4/5ths)	400 MB (500 Megabytes * 4/5ths)	14400 (4 hours)
Wednesday	0.2 (1/5th)	100 MB (500 Megabytes * 1/5th)	3600 (1 hour)

Since it is common for a download or a stream to cross the midnight boundary, most customers will usually see a fractional number of requests on the individual daily totals.

Filtering of Late-Arriving Data

Data filtering is enabled only for the Billing report and is done in order to freeze the previous month's accounting numbers.

The previous month's data received after the second day of the current month is not processed, so Reporting and RLDP numbers are slightly higher for a month due to the late arriving data. Reasons for late-arriving data include hardware maintenance, PoP connectivity issues, and hardware and software failures.

Report Data Collection Intervals

The *Limelight Control* reporting system acquires report data at different intervals depending on the data type and source.

Real-Time Delivery Reports

Other Delivery Reports

Network Transit Reports

Real-Time Delivery Reports

Data for the *Realtime Streaming* report is gathered in 5-minute increments, and is updated every 5 minutes.

Data for the *Live Stats* and *EdgeFunctions Live Stats* reports is gathered in 1-minute increments, and is updated every 30 seconds.

Other Delivery Reports

Data for non-realtime *Traffic* and *Content* reports is gathered from logs provided by Edge Servers at regular intervals. Delays from 1-20 hours can occur for specific log files depending on server load and maintenance status.

The non-realtime *Traffic* reports are:

- *Traffic*
- *EdgeFunctions Traffic*
- *EdgeFunctions Status Codes*
- *Live Push*
- *Billing*
- *DNS Overview*

The Content reports are:

- *URL Prefixes*
- *Status Codes*

Network Transit Report

Data for the *Transit* report is directly gathered from switch or router-port counters via SNMP at 5-minute intervals.

Controlling Displayed Data

A variety of controls are provided above each report to help you view exactly the data you are interested in.

Please note that report data is updated each time you change a control setting; you will need to wait for each update to complete before making additional changes.

Notes:

- Services are reported and billed based on usage
- Some services deliver content, while others provide content storage or access to Limelight's delivery backbone. Each service has its own unique characteristics and metrics.
- Not all services are applicable to every report - services are only visible in reports that make sense for the type of service

Selecting Date Ranges and Time Zones

During any Control session, if you make date range or time zone selections in any given report, Control preserves and applies your selections as follows.

- Control preserves your selections if you navigate away from the report then navigate back to it.
- Control automatically applies the selections to all other reports that you view.

Preserving and applying selections applies to these reports:

- Traffic
- Transit
- The EdgeFunctions Traffic, Status Codes and Live Stats reports
- Status Codes
- Live Push
- Origin Storage
- DNS Overview
- Realtime Streaming
- Live Stats

Note: If you log out of Control then log back in, Control does not preserve any selections you made from the previous session.

Selecting a Date Range

To control the timeframe for a report, click in the date picker. Then, in the pop-up menu, you can set the report date range by selecting a named range:

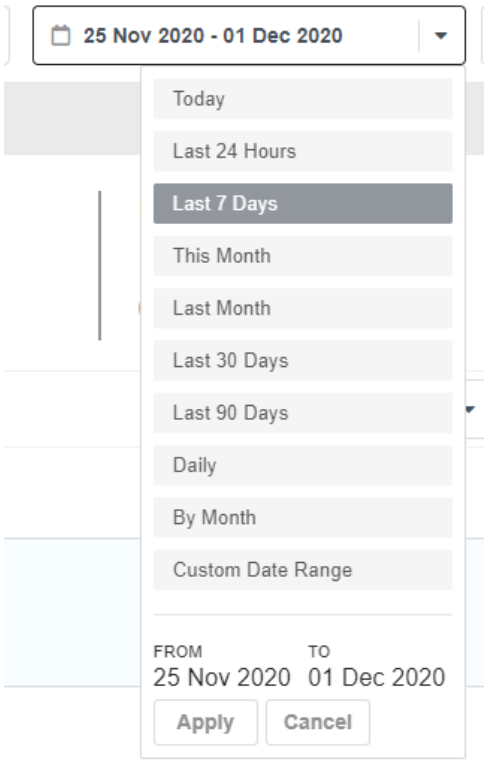

Some reports such as the URL Prefixes Report and Query String Report have limited date range selections:

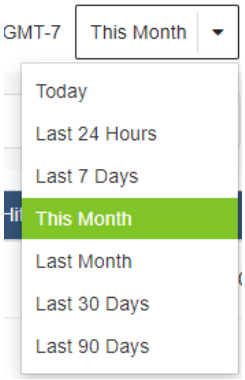

You can also pick specific dates from the Custom Date Range entry:

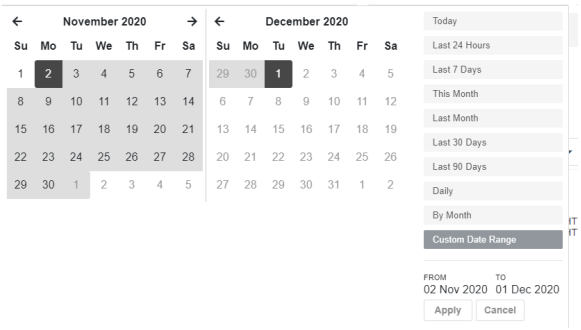

Selecting a Report Time Zone

The default time zone for all reporting data is GMT -7. You can change the time zone to reflect your local time zone or another time zone as desired.

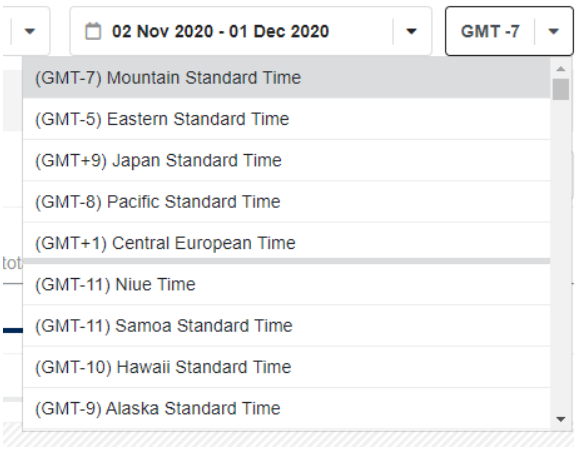

After you change the timezone, the report data updates automatically.

Viewing Data for Specific Points in a Chart

You can mouse over any data point on a chart to see the value(s) associated with that point. As you move your mouse, a vertical line is displayed to help you orient against the time axis, and the data is displayed in a popup next to your mouse.

Working with Recurring Report Emails

The Traffic Report and Status Codes Report both allow you to email report data at a specific interval and specific time of day.

Email Content

Emails consist of a message, a summary area with various statistics, and a chart. Following is an email for the Traffic Report:

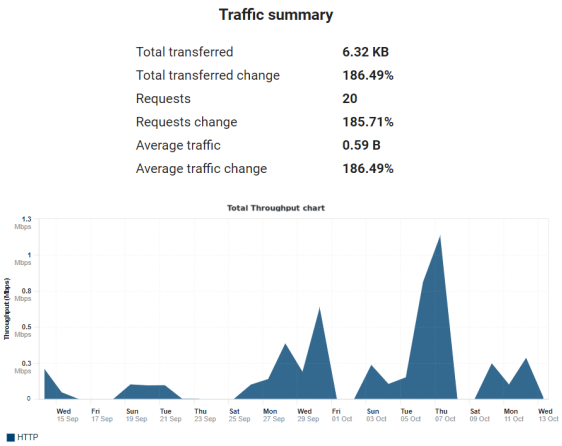

Note: The metric displayed in the email's chart (in this case 'Total Throughput') reflects the selected metric on the Overview tab.

Creating Recurring Report Emails

To create a recurring report email:

1. Click the envelope icon: on the left above the chart on the Overview tab of the Status Codes Report or Traffic Report. The *RECURRING EMAIL* dialog is displayed. The following figure shows the dialog from the Status Codes Report.

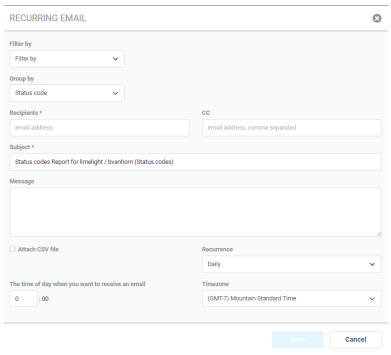

2. Fill in the fields in the dialog (see Fields in the 'RECURRING EMAIL' Dialog).
3. Click the **Save** button.

Note: After recurring report emails have been created, they are no longer visible on the Status Codes or Traffic Report pages. To edit or delete recurring report emails, you must use the My Account page. See Editing and Deleting Recurring Report Emails.

Editing and Deleting Recurring Report Emails

To edit a recurring report email:

1. Click the **Profile** icon at the top right of the Control user interface.
2. Select **My Account** from the subsequent drop-down menu.
   The **Edit My Account** page is displayed.
3. Click the **Recurring Emails** tab. For further instructions, see Managing Recurring Emails.

Fields in the 'RECURRING EMAIL' Dialog

Required fields are marked with an asterisk in the dialog.

Field	Description / Instructions
Filter by Group by	See RECURRING EMAIL Dialog - Filtering and Grouping.
Recipients	Primary notification recipient. Enter a single primary recipient email.
CC	Stands for "carbon copy." Additional email recipients. Enter one email or multiple emails separated by commas.
Subject	Email subject. Defaults to text based on the report. Modify as needed.
Message	Email body text. Enter the text for the email body.
Attach CSV file	The option to cause Control to add a CSV file of information in the report. The information is the same that you can export from the report. When you select this option, text indicating the maximum file size is displayed below the field.
Recurrence	The frequency to send the report. Select a value.
The time of day when you want to receive an email	Clock time to send the email. Use the spinner to select the hour or type a value between 0 and 23.

Field	Description / Instructions
	Note: When the dialog is opened from the Status Codes Report **URLs tab** and Traffic Report **Hosts & URLs** tab, this field is disabled and automatically set to 0700 GMT-7.
Timezone	Timezone of the time of day you configured in the **The time of day when you want to receive an email** field. Defaults to the selection in the timezone drop-down menu at the top right of the report page. You can choose a different timezone. Select a value. **Note:** The timezone field is disabled: • In the **RECURRING EMAIL** dialog for the Status Codes Report **URLs tab** and Traffic Report **Hosts & URLs** tab. • If the timezone drop-down menu at the top right of the report page is disabled for all other tabs in the Status Codes Report and Traffic Report.

RECURRING EMAIL Dialog - Filtering and Grouping

The **Filter by** and **Group by** fields allow you to control the metrics and chart groupings in the email. Filtering and grouping in the dialog are functionally equivalent to the same capabilities in the **Overview** tab's Chart. The dialog's **Filter by** and **Group by** fields are a convenience; they allow you to configure the report independently of the selections in the **Overview** tab.

Filtering - Traffic Report

Select one or at most two entries in the **Filter by**drop-down menu. Each selection causes an additional drop-down menu to display in which you can further refine your selection. For example, if you choose **Segment**, a drop-down menu of segment names is displayed.

Filtering - Status Codes Report

You can define the metrics to display in the report.

Select one or at most two entries in the **Filter by**drop-down menu. Each selection causes a control to display in which you can further refine your selection.

Selection	Additional Control / Instructions
Status codes	A field with default status codes. • Delete an existing entry by clicking the entry's **x** control. • Add a single status code by typing it and pressing Enter on your keyboard. • Add a range by typing this pattern: <start code>-<end code>

Selection	Additional Control / Instructions
	and pressing Enter on your keyboard. For example to configure the range from 200 to 203, enter: **200-203**.
Services	Drop-Down Menus with services, cache codes, or request/response types. Make a selection, then click outside the drop-down menu.
Cache codes	
Request/Response types	

Grouping - Traffic Report

By default, the chart's data grouping is by Protocol, but you can override the default by making a selection in the **Group by**drop-down menu.

Grouping - Status Codes Report

By default, the chart's data grouping is by Status code, but you can override the default by making a selection in the **Group by**drop-down menu.

Working with Email Alerts

The Traffic Report and Status Codes Report both allow you to send email alerts when a condition exceeds a threshold.

Traffic Report - For example, you might be interested in knowing if the total number of requests in the previous hour exceeds 6000. You can do so for any of your accounts and any protocols (HTTP, HTTPS, HLS, HDS). You can optionally configure one or multiple CC emails.

Status Codes Report - For example, you might be interested in knowing each time that missing file status codes (404, 410) are returned. In the Status Codes report you can create an alert from the **Overview** or **URLs** tab.

Email Content

Emails consist of a message, the date and time the alert was created, the condition, and the value that triggered the email. Following is an email for the Traffic Report.

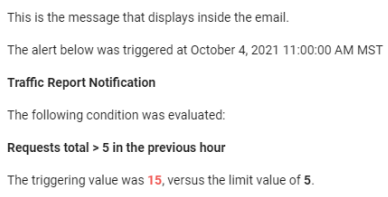

When the condition drops below the threshold value, Control sends a second alert indicating the condition is no longer met. Following is an email for the Traffic Report.

This is the message that displays inside the email.

The alert below was triggered at October 5, 2021 8:00:00 PM MST

Traffic Report Notification

The following condition is no longer met:

Requests total > 5 in the previous hour

The triggering value was **2**, versus the limit value of **5**.

Alert Icon

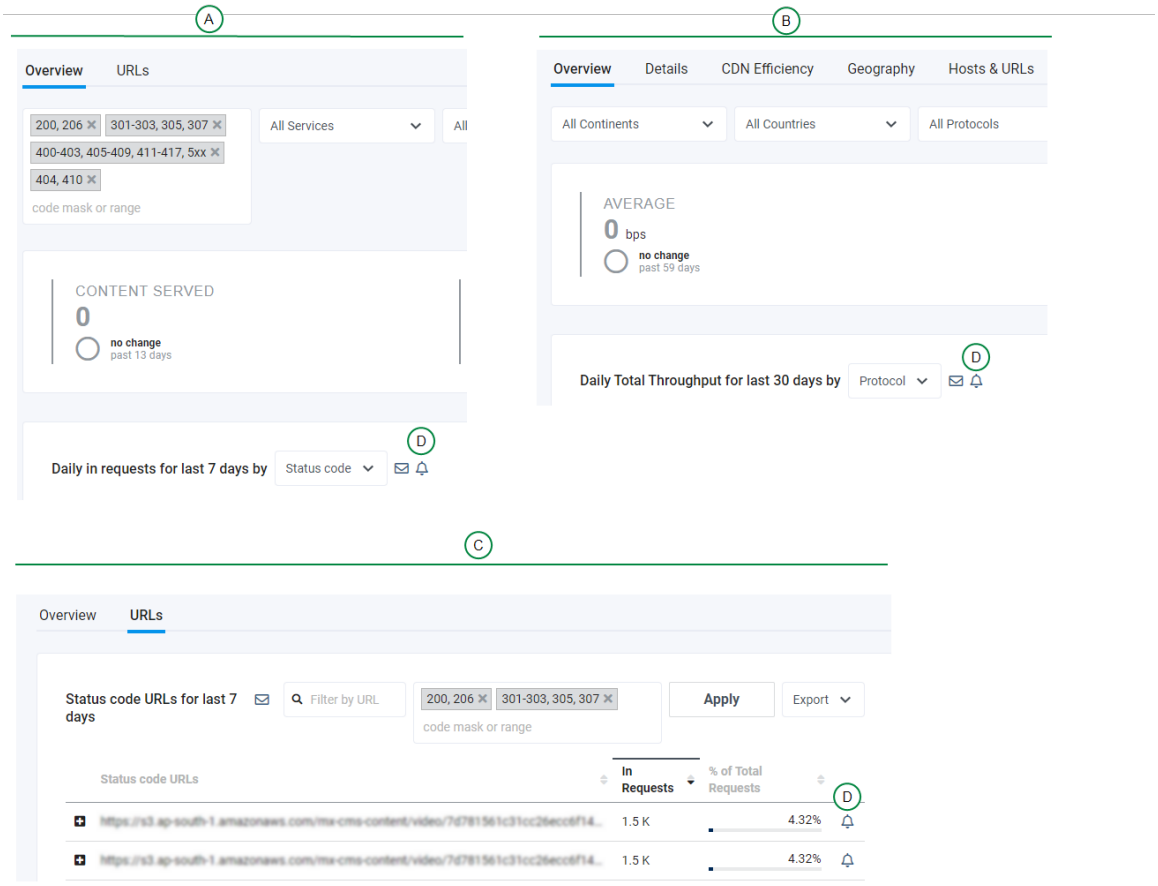

In the preceding figure:

A - Status Codes Report Overview Tab

B - Traffic Report Overview Tab

C - Status Codes Report URLs Tab

D - Bell Icon

Creating Email Alerts

To create an email alert:

1. Click the bell ico.
   The 'REPORT ALERT' dialog displays.
2. Fill out the fields in the dialog. See Fields in the 'REPORT ALERT' Dialog for details.
3. Click **Save** to save the configuration.

> **Note:** After email alerts have been created, they are no longer visible on the Status Codes Report page. To edit or delete email alerts, you must use the My Account page. See Editing and Deleting Email Alerts.

Editing and Deleting Email Alerts

To edit or delete an email alert:

1. Click the **Profile** icon at the top right of the Control user interface.
2. Select **My Account** from the subsequent drop-down menu.
   The **Edit My Account** page is displayed.
3. Click the **Alerts** tab. For further instructions, see Managing Alerts.

Fields in the 'REPORT ALERT' Dialog

Required fields are marked with an asterisk in the 'REPORT ALERT' Dialog.

Field	Description / Instructions
Accounts	A list of all your company's accounts that have the 'Alerts' product enabled. The list contains only accounts for which your user has permissions to view the Status Codes Report. An alert will be triggered only for the selected accounts. Select specific accounts or use the 'Select All' option to select all accounts.
Protocols	A list of all protocols that your company supports. An alert will be triggered only for the selected protocols. Select at least one protocol. **Note:** This field is not available for Status Codes alerts.
Recipient	Primary notification recipient. Enter a single primary recipient email.
CC	Stands for "carbon copy." Additional email recipients. Enter one email or multiple emails separated by commas.
Subject	Email subject.

Field	Description / Instructions
	Enter the text for the email's subject line. The field is automatically populated based on the account you selected in the Status Codes report.
Message	Email body text. Enter the text for the email body.
URL	The specific URL you chose to configure an alert. Read-only. **Note:** This field is only available for alerts created from the **URLs** tab of the Status Codes Report.
Send Status Code Report when, Send Traffic Report when	Conditions under which you want to email to be sent. See Conditions for Sending an Email.

Conditions for Sending an Email

Use the 'Send Status Codes Report when' or 'Send Traffic Report when' section to define the threshold for sending an email alert.

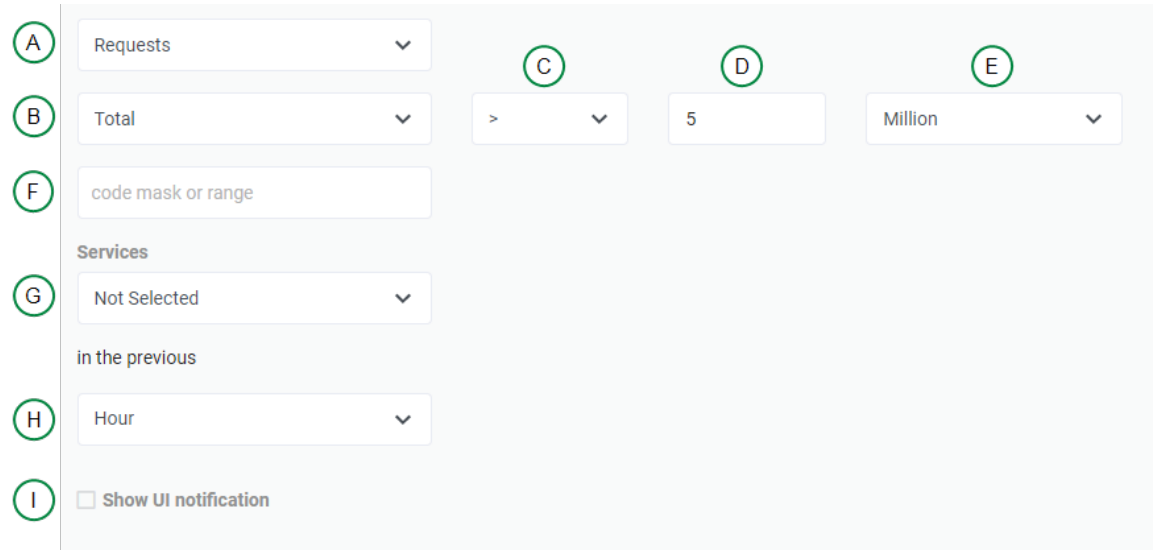

The fields identified in the previous screenshot are described in the following table. For sample usages, see Example 1 and Example 2.

Field	Description
A - Threshold Metric	Metric associated with threshold: 'Requests', 'Requests In', and so on.
B - Total or Percent Change	Means of expressing the change of the selected metric. 'Total' - change expressed as a sum of the selected metric over the time unit selected in the 'in the previous' drop-down menu. 'Percent Change' - change expressed as a percentage from the time unit in the 'in the previous' drop-down menu.
C - Greater Than or Less Than	Inequality operator to indicate the threshold comparison value.
D - Threshold Number	Value beyond which an email is triggered. **Note:** If you select 'Percent Change' in the 'Total or Percent Change' field, this field represents a percentage.
E - Threshold Numeric Unit	Unit of the selected metric in the 'Threshold Metric' drop-down menu: 'Thousand', 'Million', and so on. **Notes:** • This field is visible only if you select 'Total' in the 'Total or Percent Change' field. • Use the first entry (-) to indicate the entry in the 'Count or Percentage Change' value is a straight count.
F - Code Mask or Range	HTTP response code(s) associated with the threshold. **Note:** This field is not available for Traffic Report alerts.
G - Services	Protocols associated with the threshold. **Note:** This field is not available for Traffic Report alerts.
H - in the previous	Time unit for the threshold: '5 Minutes', 'Hour', 'Day'.

Field	Description
I - Show UI notification	If this field is checked, then when the threshold has been crossed, Limelight Control displays an alert popup in the user interface in addition to sending an email.

Example 1

If you want to receive an alert when the number of incoming requests with a 202 or 206 status for the HLS service is 500 more from the previous day, make the following selections:

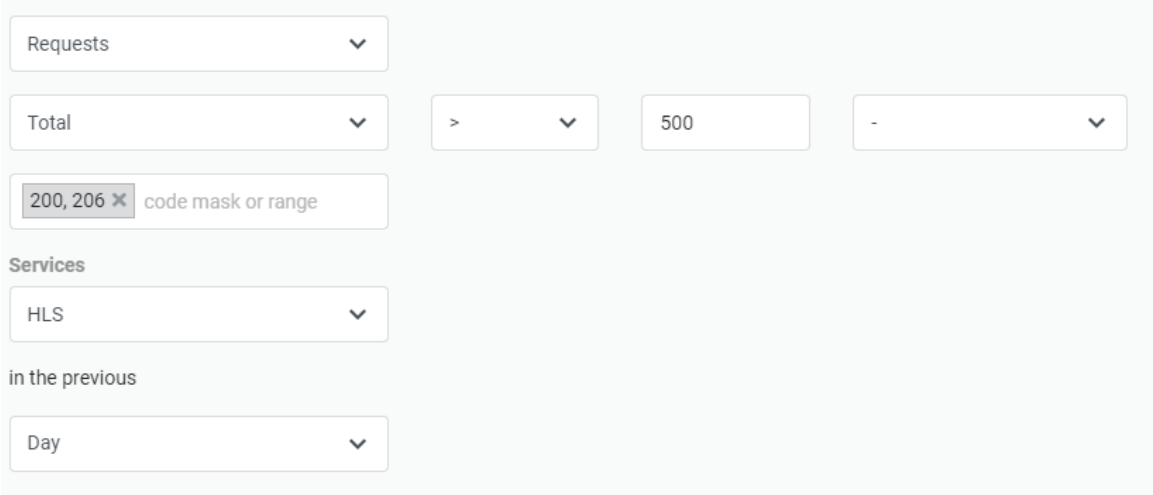

Example 2

If you want to receive an alert under the same conditions as Example 1, but express the change as a percentage from the previous day, make the following selections:

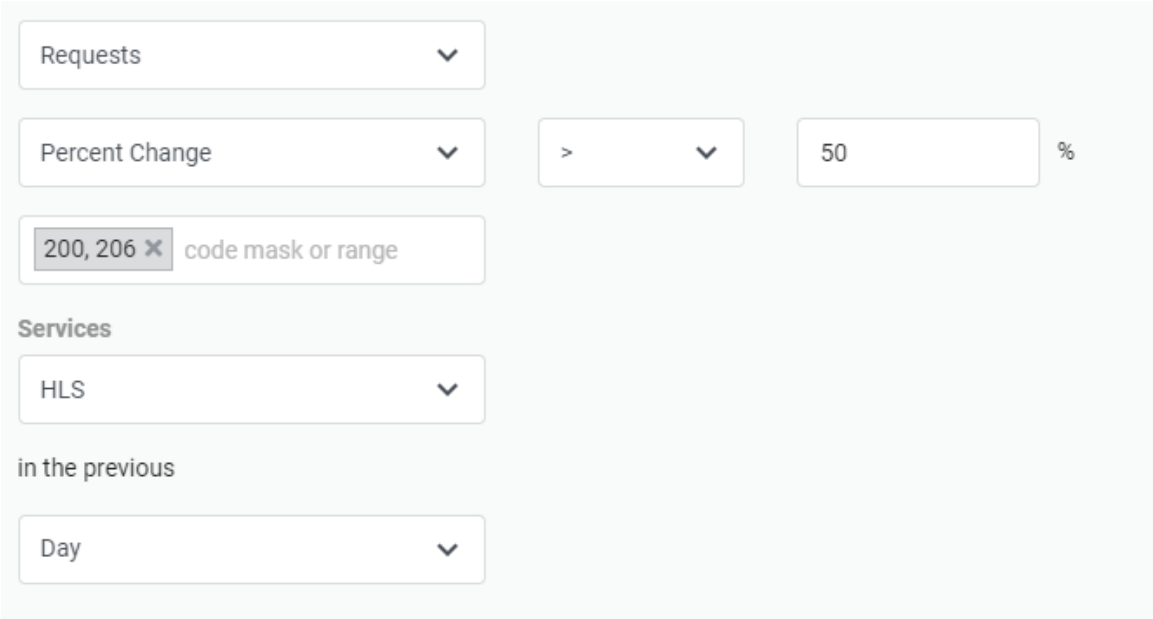

Traffic Reports

Traffic reports display summary data about CDN usage, such as data volume and the number of associated requests and connections, by service (protocol). Traffic reports can include both high-level overviews and detailed data.

Billing Report

The *Billing Report* shows monthly traffic for your active products and services and your Limelight accounts. You can use data in the report for various reasons such as cost estimates for future Limelight invoices.

> **Note:** Due to data filtering and proration, data in the Billing report may not align with data in other individual reports. For more information, see Filtering of Late-Arriving Data.

Billing Types

The data in the report depends on the selected billing type.

Billing Type	Metrics	Dimension
Content Delivery	95th Percentile of Inbound (Mbps), 95th Percentile Of Outbound (Mbps), 95th percentile of the total (Mbps) Data Transfer in (GB), Data Transfer out (GB), Data Transfer Total (GB) Requests (10Ks)	Service or account, depending on the selection in the drop-down menu above the table on the right.
Discrete Storage	Peak Storage (GB)	Account
DNS	Number of Outbound Requests	Account
DRM	Number of Outbound Requests	Account
EdgeFunctions	Count of Invocations, Compute Usage (allocated memory x function execution time)	Account
Intelligent Ingest	Total GB Transferred, Number of Sessions (Requests)	Account
Origin Storage	Count of Unique Objects in, Count of Unique Bytes in Origin Storage (GB), Total Bytes Retrieved from Origin Storage (GB)	Policy or account, depending on the selection in the drop-down menu above the table on the right.
Live Push Ingest	Total Ingest Traffic (GB)	Account
MMD Live	*Data Transfer IN, Transcode bytes TOTAL, Transcode pixels TOTAL	Account

Billing Type	Metrics	Dimension
MMD Live to VOD	Number of Minutes Recorded	Account
Realtime Streaming	Peak Ingest (Kbps), Data Transferred in (bytes), Viewer Hours	Account
SSL SNI Certificate Hosting	Count of published certificates	Account
Transit	Data Transfer in (GB), Data Transfer out (GB), Data Transfer Total (GB) 95th Percentile of Inbound (Mbps), 95th Percentile of Outbound (Mbps), 95th percentile of the total (Mbps), 95th Percentile High of Inbound and Outbound Mbps	Account

* Transcode bytes TOTAL is the total number of bytes transcoded. Transcode pixels TOTAL is the total number of pixels transcoded during the process of transcoding an RTMP stream. Transcoding is the process of receiving an RTMP stream with a single bitrate and outputting it to multiple bitrates.

Understanding The Report

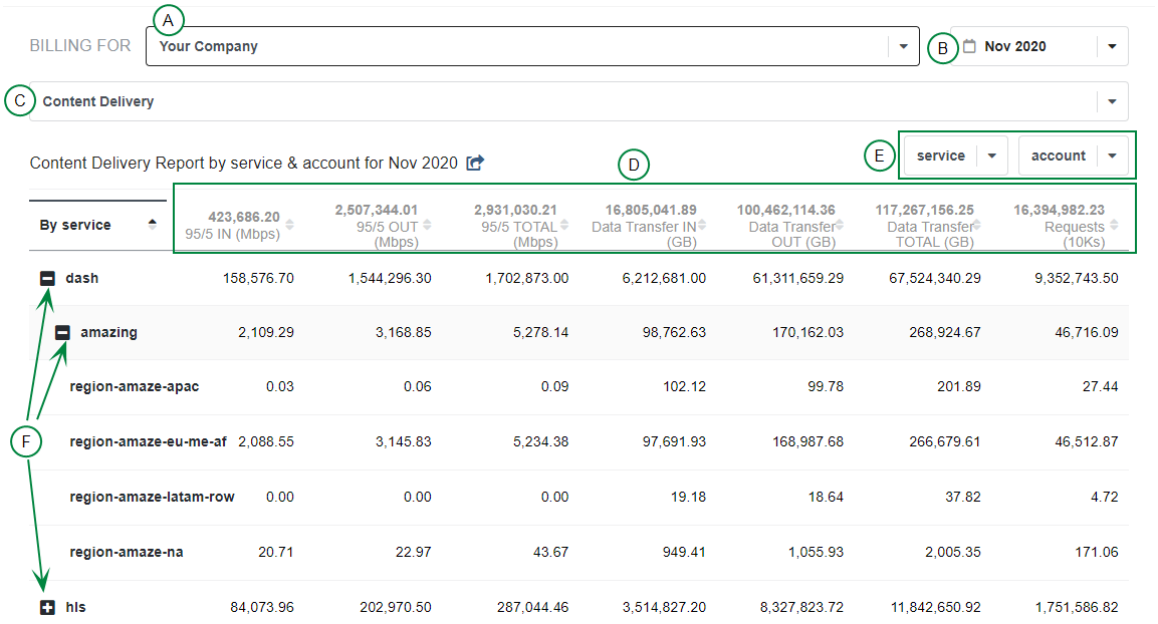

The following list identifies controls in the report:

A - Parent company and child companies drop-down menu - if the account you selected for your Control session has child companies, the parent company and its child companies are displayed in a drop-down menu. Otherwise, the parent company is displayed in a text field. See Selecting Child Companies for additional information.

B - Reporting month drop-down menu - Reporting Period.

C - Billing type drop-down menu - determines data to appear in the report.

D - Report data - Data in the report.

E - Groupings drop-down menus - Control hierarchies and expandable sections.

F - Expandable selections - provide visibility into lower-level groupings.

You can select accounts, services, reporting month, and up to two levels of groupings.

> **Note:** Letters correspond to call-outs in the preceding screenshot and do not imply sequencing.

Selecting Child Companies

If the account you selected has child companies, the parent company and its child companies are displayed in a drop-down menu at the top of the page.

To limit report data to one or more child companies, select the desired child companies in the **Parent company and child companies drop-down menu**, then click the **apply** button at the bottom of the list.

If you don't select a child company, data for the parent company is displayed.

> **Note:** The drop-down menu is available only if the account has child companies.

Interactions Between Grouping Controls and Report

The left-most column in the table has expandable selections; the data displayed is hierarchical and depends on the selections you make in the Groupings drop-down menus.

The number of Groupings drop-down menus and their content depends on the account and services you select. (Services available depend on your active Limelight services.)

Values in the Groupings drop-down menus are 'Account', 'Service', 'Region'. Depending on your services and accounts, one, two, or no Groupings drop-down menus may be displayed:

- If just two values are available, you will see only one drop-down menu. The table automatically shows the hierarchy for the selection you make.
  For example, if the drop-down menu contains only 'Service' and 'Account', and you select 'Account', all accounts are listed in the table. Expand any account to see the services under the account.
- If three values are available, you will see two drop-down menus. The first provides the major grouping. The second shows the groupings not selected in the first drop-down menu.
  For example, If the available values are 'Account', 'Service', 'Region', and you select 'Account' in the first drop-down menu, the second drop-down menu will contain 'Region' and 'Account'.
- No drop-down menus appear if only one value is available.
  For example, if only 'Account' is available, the accounts are listed in the report table.

Examples of Expanding Groups

If you select 'Service' and 'Account' respectively in the Groupings drop-down menus, the first-level groups are your services. Expanding each of them reveals groups of the accounts you selected.

As another example, assume you have two accounts and have the 'http' and 'https' services and operate in four regions. If you select 'Account' in the first drop-down menu, the second drop-down menu's remaining choices are 'Service' and 'Region'. If you then select Service, you will see hierarchical entries in the first column that might look something like this when expanded:

```
first account
  http
    first region
    second region
second account
  https
    third region
    fourth region
```

Working with Report Data

Making Selections

You can select accounts, reporting month, billing type, and up to two levels of groupings.

1. Select a service in the service drop-down menu.
2. Select one or more accounts in the account drop-down menu.
3. Select a month in the Reporting month drop-down menu.
4. Select a service from the Billing type drop-down menu.
5. If Groupings drop-down menus are available, make the desired selections.

After you make selections, the report title, and content change to reflect your selections.

Removing and Adding Columns

Remove a column by hovering your mouse pointer over the column header and clicking the **x** icon on the upper right of the column header. Doing so removes the column immediately.

If there are additional columns to display, a **+** icon displays on the right side of the table header row. Click the icon and choose a column to add.

Sorting the Report

By default, the report is sorted in ascending order by the data in the first column. Click the 'Sort Direction' icon to sort data in descending order.

Exporting Report Data

You can export report data to a Comma-Separated Values (CSV) file.

Click the **Export** drop-down menu on the right above the table, then choose a granularity:

Granularity	Aggregation/Grouping in the CSV File
Monthly	Data is aggregated by month.

Granularity	Aggregation/Grouping in the CSV File
Daily	Data is broken out by day for the month selected in the Reporting month drop-down menu.

After you choose a granularity, Control creates and downloads the report.

> **Note:** The SSL SNI reporting type does not have a Daily option because SSL SNI does not support daily granularity.

DNS Overview Report

The DNS Overview report allows you to view DNS Traffic over time. You can use this report to track the number of requests for specific hostnames and accounts.

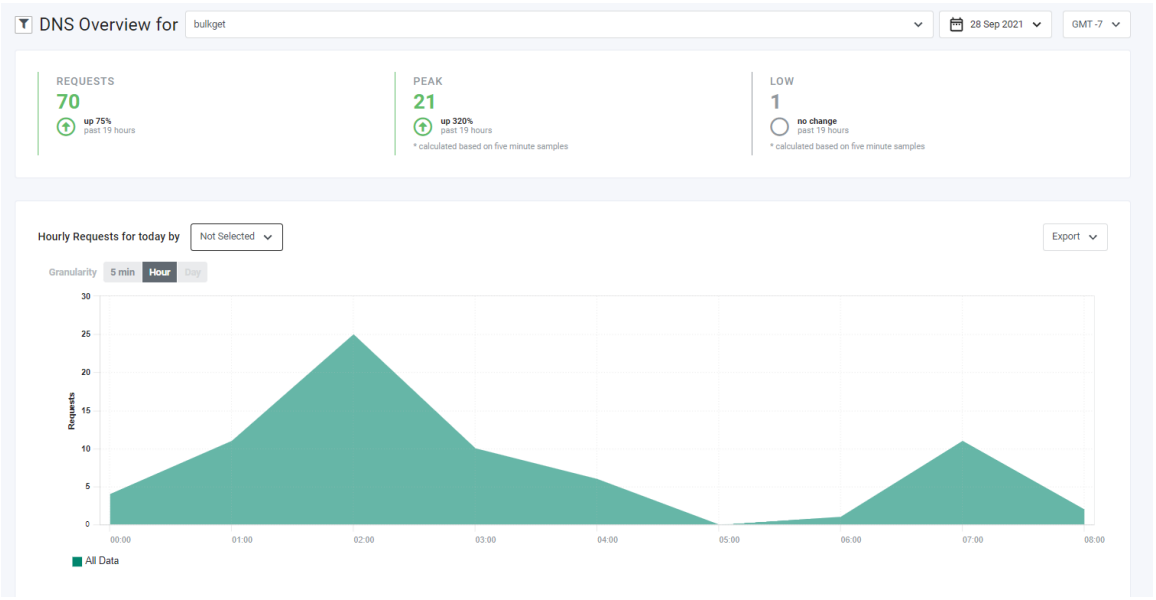

The report consists of a summary area and an aggregate data line chart.

Report Specifications

Latency	Latency depends on the selected granularity. **Granularity** / **Latency** 5 Minutes / 5 to 10 minutes Hour / 1 hour + delta* Day / 1 day + delta* *delta = approximately 5-10 minutes
Granularity	5 minutes, hour, day
Dimensions	• Account • Hostname
Metrics - Summary Area	• Requests • Peak • Low
Metrics - Chart	Requests

Delivery Mechanism	Realtime Reporting API
Associated API Endpoint (s)	`/realtime-reporting-api/dns/` • GET - Retrieves DNS report data based on query parameters • POST - Retrieves DNS report data based on the request body

Selecting Accounts, Date Range, and Time Zones

You can make selections in the controls above the tab header:

- **DNS OVERVIEW FOR**. Select one or more accounts to which you have access for cross-account analysis. Click the **Select All** button to select all accounts.
- **Date range**. Pick from pre-set time frames or choose custom date ranges in the drop-down menu. Click the **Apply** button on custom ranges.
- **Time zone**. The top five most commonly used timezones in Control are at the top of the drop-down menu. Scroll down for additional time zones.

Summary Area

For various statistics, the Summary Area shows the percent change for the selected reporting date range relative to the previous time range of the same duration. Colors and arrows represent changes:

- An increase displays in green with an arrow pointing up.
- A decrease displays in red with an arrow pointing down.
- If there was no change from the previous period, the text is gray with an empty circle instead of an arrow.
- Percentage up or down is in bold text to the right of the arrow or circle.
- Information in gray text beneath percentage up or down indicates the total date range covered minus any time remaining in the current period. For example, if you selected **Last 7 days** as the reporting period and the current date is part of the past seven days, the information displayed is for the past 13 days.

Statistics in the Summary Area are:

Statistic	Description
Requests	Total requests for the selected accounts, time period, and timezone.
Peak	Highest number of requests for the selected accounts, time period, and timezone.
Low	Lowest number of requests for the selected accounts, time period, and timezone.

Filtering by Hostname

The filter icon (a funnel) on the left side of the tab header allows you to filter by hostname. The icon is a toggle that displays or hides the filter control.

1. Click the filter icon.
2. Make a selection in the subsequent **Filter by** drop-down menu.
   By default, the **Hostname** entry is unchecked. As such, data for all hosts is displayed in the chart.
3. If desired, select **Hostname** to display a drop-down menu with hostnames you can select for display in the chart.
4. Click the **Apply** button.

Selecting a Data Grouping

On the left beneath the Summary area is the **Grouping** drop-down menu that allows you to determine how data is broken out in the chart: by Account or Hostname.

Make a selection.

Selection	Data Displayed in Chart	Information Displayed in Legend
Not Selected	All data for the selected metric (See Choosing Metrics.)	Single entry: **All Data**
Account	All accounts that you selected at the top of the page	A list of the selected accounts.
Hostname	All hostnames you selected for filtering	A list of the filtered hostnames.

The chart reflects your selection.

Choosing Granularity

The chart can be further refined by selecting one of the **Increment** values:

- 5 min
- Hour
- Day

The selection you make determines increments along the Y-axis.

Exporting Chart Data

To export data currently displayed in the chart, click the **Export** drop-down menu on the right above the chart; then choose an option:

- **PowerPoint**: Export to a PowerPoint file that contains a screenshot of the chart.
- **CSV**: Export data to a Comma-Separated Values (CSV) file that reflects the currently selected chart options.

After you choose an option, Control creates and downloads the report.

Toggling Chart Data

A legend below the chart identifies the chart content by color. The legend reflects the choices that you made in the **Hostname Filter** and the **Grouping**drop-down menu.

For example, if you filtered by two specific hostnames and selected **Hostname** in the drop-down menu, data for the hostnames displays in the chart. Labels for the two names display in the chart legend.

The legend labels are toggles that you can use to display or hide the corresponding chart data. By default, all labels are toggled on. Click a label to hide or show the corresponding data in the chart. Labels that are toggled off have a gray font color.

How Metrics Are Calculated

Summary Area

Metric	Calculation
Requests	No calculation. Data provided by EdgeQuery.
Peak	No calculation. Data provided by EdgeQuery.
Low	No calculation. Data provided by EdgeQuery.

Chart

Metric	Calculation
Requests	No calculation. Data provided by EdgeQuery.

EdgeFunctions Live Stats Report

The EdgeFunctions Live Stats Report shows EdgeFunctions data and can be used to monitor function characteristics such as efficiency and usage. The report displays data for the past 12 hours. Data is automatically refreshed every 30 seconds.

> **Note:** To disable or enable auto-refreshing of chart data, click the **Auto refresh** toggle on the right side of the page underneath the **Time Zone Control.**

The report consists of a summary area and an aggregate data line chart.

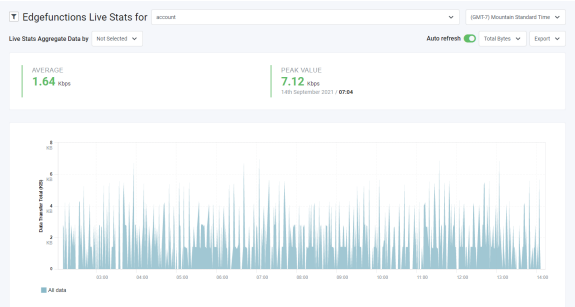

Report Specifications

Latency	15 minutes
Granularity	One minute (cannot be modified by customers)
Dimensions	• Account • Shortname • Function
Metrics - Summary Area	Average value of the chosen metric. Peak Value and day/time value occurred for the chosen metric. For details, see How Metrics Are Calculated.
Metrics - Chart	• In bytes, Out bytes, Total bytes (in + out) • Execution time • Duration • Requests • Compute usage For details, see How Metrics Are Calculated.
Delivery Mechanism	Realtime Reporting API
Associated API Endpoint(s)	/realtime-api/edgefunctions/livestats

Choosing a Time Zone

Use the **Time Zone Control** at the top right of the page to select a timezone. The top five most commonly used timezones in Control are at the top of the drop-down menu. Scroll down for additional time zones.

Choosing Accounts

You can choose accounts for which you wish to see EdgeFunctions Live Stats. Click the drop-down menu to the right of the *EDGEFUNCTIONS LIVE STATS FOR* title and select the desired accounts. The chart data is limited to the accounts you select.

> **Note:** You must select at least one account; otherwise, the default company is automatically selected, and a warning is displayed.

Choosing Functions to Display

Use the **Function Filter** to determine which functions to display in the chart.

The filter icon (a funnel) on the left side of the *EDGEFUNCTIONS LIVE STATS FOR* title allows you to filter by function name. The icon is a toggle that displays or hides the function filter control.

1. Click the filter icon.
2. Make a selection in the subsequent **Filter by**drop-down menu.
   By default, the **Function** entry is unchecked. As such, data for all functions is displayed in the chart.
3. If desired, select **Function** to display a drop-down menu with function names you can select for display in the chart.
4. Click the **Apply** button.

Choosing Data Breakout in the Chart

Beneath the *EDGEFUNCTIONS LIVE STATS FOR* title is the **Live Stats Aggregate Data by**drop-down menu that allows you to determine how data is broken out in the chart: by account, function, or both accounts and function.

Make a selection.

Selection	Data Displayed in Chart	Data Displayed in Legend
Not Selected	Accounts and functions.	Single entry: **All Data**.
Account	Accounts only.	A label for each account name.
Function	Functions only.	A label for each function name.

The chart reflects your selection, and the Summary Area displays the average and peak value for the selected metric.

Choosing Metrics

Beneath the **Timezone Control** is a drop-down menu that allows you to select chart and Summary Area metrics.

Make a selection.

Selection	Description
Total Bytes	Total of In Bytes + Out Bytes, expressed in KB.
In Bytes	Bytes ingressed by the function, from all sources, expressed in KB.
Out Bytes	Bytes egressed by the function to all destinations, expressed in KB.
Execution Time	Total time all functions took to run, expressed in seconds.
Compute Usage	Allocated memory x function execution time.

The chart and chart legend reflect your selection.

Toggling Chart Data

A legend beneath the chart identifies the chart content by color. The legend reflects the choices that you made in the **Function Filter** and the **Live Stats Aggregate Data by**drop-down menu.

For example, if you filtered by two specific function names and you selected **Function** in the drop-down menu, data for the functions displays in the chart. Labels for the two names display in the chart legend.

The legend labels are toggles that you can use to display or hide the corresponding chart data. By default, all labels are toggled on. Click a label to hide or show the corresponding data in the chart. Labels that are toggled off have a gray font color.

Exporting Chart Data

To export data currently displayed in the chart, click the **Export** drop-down menu on the right above the chart; then choose an option:

- **PowerPoint**: Export to a PowerPoint file that contains a screenshot of the chart.
- **CSV**: Export data to a Comma-Separated Values (CSV) file that reflects the currently selected chart options.

After you choose an option, Control creates and downloads the report.

How Metrics Are Calculated

The Summary Area and chart show values for the chosen metric for the functions selected in the **Function Filter**. The unit of measurement depends on the chosen metric.

Metrics are Total Bytes, In Bytes, Out Bytes, Execution Time, Requests, Compute Usage.

Summary Area

The Summary Area shows the Average and Peak Value. The peak value includes the day/time the value occurred.

Metric	Calculation
Average for the chosen metric	Sum of all measurements for the metric across time range, divided by 12.
Peak Value for chosen metric.	Largest measurement for the metric over the last 12 hours.

Chart

Metric	Calculation
Total Bytes	Sum of In bytes and Out bytes.
In Bytes	Sum of bytes ingressed by the functions from all sources.
Out Bytes	Sum of bytes egressed by the functions to all destinations.
Execution Time	Sum of function run times.
Requests	Sum of function invocations.
Compute Usage	Product of Allocated memory and function execution time.

EdgeFunctions Status Code Report

The *EdgeFunctions Status Codes Report* provides a summary of various status code counts for a given account, reporting date range, and timezone.

The report also shows the number of invocations (requests) per status codes, functions, or accounts.

The report consists of a *Summary Area* and a chart.

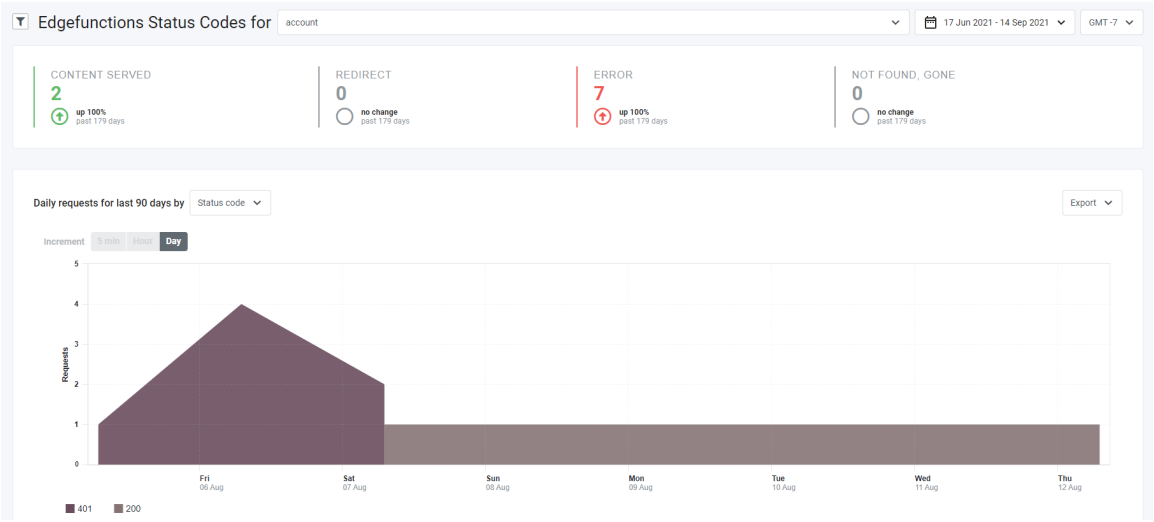

Selecting Accounts, Date Range, and Time Zones

You can make selections in the controls above the *Summary Area*.

- *EDGEFUNCTIONS STATUS CODES FOR*: Select one or more accounts to which you have access for cross-account analysis. Click the **Select All** button to select all accounts.

  > **Note:** You must select at least one account; otherwise, the default company is automatically selected, and a warning is displayed.

- **Date range**. Pick from pre-set time frames or choose custom date ranges in the drop-down menu. Click the **Apply** button on custom ranges.
- **Time zone**. The top five most commonly used timezones in Control are at the top of the drop-down menu. Scroll down for additional time zones. Select a timezone.

Filtering

A filter drop-down menu directly above the *Summary Area* allows you to filter by status codes, functions, and accounts.

> **Note:** Filters you apply only affect the chart itself and not the *Summary Area*.

You can filter on at most two items, so when you choose two items, the excluded choices are disabled and cannot be selected. Each selection causes additional filters to display. You use the filters to further refine your choices. When you filter by status codes, the additional filter presents a list of default choices, which correspond to the headings in the *Summary Area*.

Selecting Two Items for Filtering

If you filter by status code and also functions or accounts, the chart reflects only functions or accounts, depending on which you selected.

If you filter by both functions and accounts, the chart reflects only functions.

Make the desired selections, then click the **Apply** button.

> **Note:** If you modify the selections in the *EDGEFUNCTIONS STATUS CODES FOR*drop-down menu, the filter selections you made are reset to defaults and the chart updates automatically to match the defaults.

Summary Area

For various statistics, the *Summary Area* shows the percent change for the selected reporting date range relative to the previous time range of the same duration. Colors and arrows represent changes:

- An increase displays in green with an arrow pointing up.
- A decrease displays in red with an arrow pointing down.
- If there was no change from the previous period, the text is gray with an empty circle instead of an arrow.
- Percentage up or down is in bold text to the right of the arrow or circle.
- Information in gray text beneath percentage up or down indicates the total date range covered minus any time remaining in the current period. For example, if you selected **Last 7 days** as the reporting period and the current date is part of the past seven days, the information displayed is for the past 13 days.

> **Note:** Information in the *Summary Area* is determined only by the selected account and reporting date range. It is not affected by the selections you make in any other controls on the page.

Statistics in the *Summary Area* are:

- *CONTENT SERVED*: content was returned to the requestor
- *REDIRECT*: requests were redirected
- *ERROR*: requests resulted in an error condition
- *NOT FOUND, GONE*: requested file not found

Available Controls

Report Granularity

Above the report are toggles for choosing granularity: 5 minutes, hour, day. Toggles are active, depending on the reporting date range.

Report Legend

A legend beneath the chart identifies the chart content by color. Available legend labels depend on the filters you selected. The legend labels are toggles that you can use to display or hide the corresponding chart data.

Interacting with the Report

Make selections in the controls on the page as needed.

Hover over the chart to view data details in tooltip popups.

> **Note:** Anytime you see a ~ in your data, it means that the number is an approximation.

To export data currently displayed in the chart, click the **Export** drop-down menu on the right above the chart; then choose an option:

- **PowerPoint**: Export to a PowerPoint file that contains a screenshot of the chart.
- **CSV**: Export data to a Comma-Separated Values (CSV) file that reflects the currently selected chart options.

After you choose an option, Control creates and downloads the report.

EdgeFunctions Traffic Report

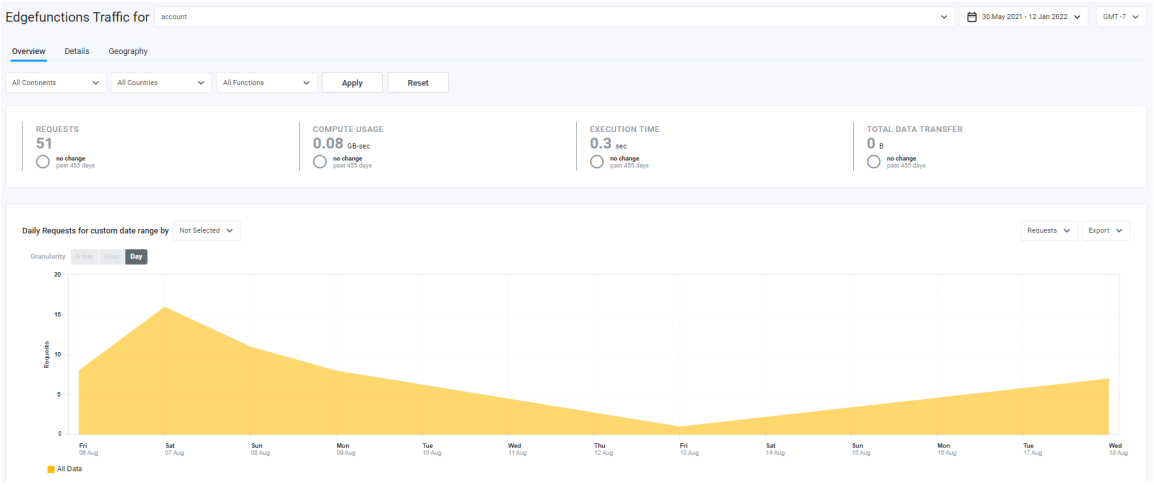

The EdgeFunctions Traffic provides a granular view of EdgeFunctions activity for Limelight's Content Delivery customers. You can view this information: number of function requests, CPU usage, and various data transfer metrics. The report has these tabs:

- Overview
- Details
- Geography

Report Specifications

Latency	Two minutes
Granularity	5 minutes, hour, day
Dimensions	Account, function name, continent, country
Metrics	Requests, compute usage, execution time, total data transfer, transfer in, transfer out See How Metrics Are Calculated for details.
Delivery Mechanism	Realtime Reporting API
Associated API Endpoints	• `/realtime-reporting-api/edgefunctions/` returns EdgeFunctions traffic data • `/realtime-reporting-api/edgefunctions/geo` returns EdgeFunctions traffic geographic data

Selecting Accounts, Date Range, and Time Zones

You can make selections in the controls above the tab header:

- *EDGEFUNCTIONS TRAFFIC FOR*: Select one or more accounts for which you have access for cross-account analysis. Click the **Select All** button to choose all accounts.
- **Date range**. Pick from pre-set time frames or choose custom date ranges in the drop-down menu. Click the **Apply** button on custom ranges.
- **Time zone**. The five most commonly used timezones in Control are at the top of the drop-down menu. Scroll down for additional time zones. Select a timezone.

Filtering Report Data

Using filter controls at the top of the report, you can filter on **Continents**, **Countries**, and **Functions**.

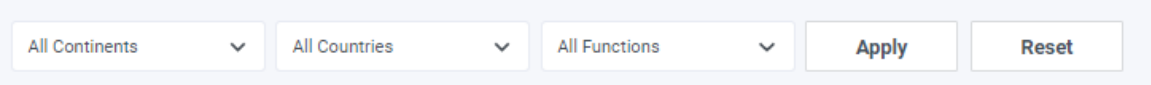

Make the desired selections; then click the **Apply** button to apply your filter choices.

To reset filters to the default, click the **Reset** button.

Notes:
- Filter controls default to **All**.
- Some filter selections are mutually exclusive. For example, if you select **North America** as the continent, you cannot select individual countries.
- The **Geography** tab only contains filter controls for **Functions**.
- Filter selections you make on one tab are automatically applied to other tabs.
- When you make filter selections on a tab then navigate to another tab, then navigate back to the original tab, the selections are preserved in the original tab.

Overview Tab

This tab presents a synopsis of EdgeFunctions data for the selected account and reporting data range.

Summary Area

For various statistics, the *Summary Area* shows the percent change for the selected reporting date range relative to the previous time range of the same duration. Colors and arrows represent changes:

- An increase displays in green with an arrow pointing up.
- A decrease displays in red with an arrow pointing down.
- If there was no change from the previous period, the text is gray with an empty circle instead of an arrow.
- Percentage up or down is in bold text to the right of the arrow or circle.
- Information in gray text beneath percentage up or down indicates the total date range covered minus any time remaining in the current period. For example, if you selected **Last 7 days** as the reporting period and the current date is part of the past seven days, the information displayed is for the past 13 days.

Note: Information in the *Summary Area* is determined only by the selected account and reporting date range. It is not affected by the selections you make in any other controls on the page.

Statistics in the *Summary Area* are:

- *REQUESTS*: Total number of times functions were called, measured in thousands.
- *COMPUTE USAGE*: (allocated memory) x (function execution time).
- *EXECUTION TIME*: Cumulative number of seconds that functions have run within the given time.

- *TOTAL DATA TRANSFERRED*: Total of all data transferred in and out as a result of function calls, measured in megabytes.

Available Controls

Beneath the *Summary Area* are drop-down menus for breaking out chart data and choosing metrics to display in the report.

- On the left, you can break out chart data: **Not Selected**, **Account**, **Function Name**, **Continent**, **Country**.

  > **Note: Not Selected** indicates to display all data rather than breaking out by the other values.

- On the right, you select metrics to display: **Requests**, **Data Transfer In**, **Transfer Out**, and **Transfer Total**, **Execution Time**, **Compute Usage**. The selection you make is reflected in the report title.

Above the report are toggles for choosing granularity: 5 minutes, hour, day. Toggles are active, depending on the reporting date range.

Interacting with the Tab

Make selections in the controls on the tab as needed.

Hover over the chart to view data details in tooltip popups.

Click labels to the right of the chart to display or hide chart data.

To export data currently displayed in the chart, click the **Export** drop-down menu on the right above the chart; then choose an option:

- **PowerPoint**: Export to a PowerPoint file that contains a screenshot of the chart.
- **CSV**: Export data to a Comma-Separated Values (CSV) file that reflects the currently selected chart options.

After you choose an option, Control creates and downloads the report.

Details Tab

This tab provides the same information as the **Overview** tab but in a tabular format. The metrics available for selection in the **Overview** tab are presented as table headings in the **Details** tab.

Available Controls

Beneath the tab header is a drop-down menu for breaking out data in the report, by **Account**, **Function Name**, **Continent**,or **Country.** The selection you make determines the first-level grouping in the table.

Interacting with the Tab

Make selections in the drop-down menus on the tab as needed.

Sort the data by clicking any column header. A black triangle to the right of the header text indicates the sort direction:

- Triangle points up: Rows are in ascending order.
- Triangle pointing down: Rows sorted in descending order.

The table does not have a secondary sort order.

Expand a row in the table by clicking the **+** icon. Doing so reveals entries with a unit of Month. Expand a month by clicking the **+** icon to reveal days in the month. Finally you can click the **+** icon in a day row to drill down to specific hours in the day.

Removing and Adding Columns

Remove columns by hovering over a column header and clicking the **x** icon on the upper right of the column header. Doing so removes the column immediately.

If there are four or fewer columns, a **+** icon displays on the right side of the table header row. Click the icon and choose a column to add.

To export table data to a Comma-Separated Values (CSV) file, click the **Export as CSV**drop-down menu on the right above the table, then choose an option.

Option	CSV File Contains
Monthly	A total for each metric in each month.
Daily	A total for each metric in each day in each month.
Hourly	A total for each metric in each hour in each day in the month.

After you choose an option, Control creates and downloads the report.

> **Note:** If the table contains no data, the CSV file contains only headings.

Geography Tab

This tab shows EdgeFunctions traffic on an interactive map. The map can be used to analyze EdgeFunctions-specific traffic on a geographic basis.

Available Controls

Drop-down menus under the tab header allow you to group data and select metrics to display in the chart.

- On the left, you can break out map data: **Not Selected**, **Account**, or **Function Name**.

  > **Note:** **Not Selected** indicates to display all data rather than breaking out by the other values.

- On the right, you can choose metrics: **Requests**, **Data Transfer In**, **Transfer Out**, **Data Transfer Total**, **Execution Time**, **Compute Usage**.

Relative traffic volume is indicated by the darkness of colors for each geographic area. Darker colors indicate higher traffic volumes.

Interacting with the Tab

Make selections in the controls on the tab as needed.

Controls for zooming allow you to focus in or out:

- Use the **[+]** and **[-]** controls on the upper left of the map to zoom in to the continent level and zoom out to the world level.
- Click a region to zoom in (continent, country, or state) and use the minus to zoom back out. You can also use the **Continent** and **Whole World** controls at the top right of the map to zoom out.
- Click a continent to view data at the country level.

How Metrics Are Calculated

Metric	Calculation
Requests	No calculation. Data provided by EdgeQuery.
Compute usage	Product of allocated memory and function execution time.
Execution time	Sum of function run times.
Data transfer in	No calculation. Data provided by EdgeQuery.
Data transfer out	No calculation. Data provided by EdgeQuery.
Total data transfer	Sum of Data transfer in and Data transfer out.

LDS Overview Report

The LDS Overview Report allows you to check Log Delivery Service (LDS) data latency, verify data completeness, and view the volume of information uploaded to log files.

Concepts

When a request for your content enters the CDN, EdgePrism logs the request. Sometime later, Log Delivery Service gathers your log entries and then delivers them in files in one of three locations, depending on your Log Delivery Service configurations:

- Limelight's Origin Storage
- The Google Cloud Platform
- Amazon S3 Buckets

For information about Log Delivery Service configurations, see Configuring Log Delivery Service.

Despite Limelight's robust network capabilities, issues may occur.

- **Latency**: There is sometimes a delay involved delay involved between the time a request is logged on edge servers and the time the request is delivered in Log Delivery Service files to the storage.
- **Data Completeness**: On occasion, only a subset of log entries will be delivered. Data completeness is a comparison of the number of log lines uploaded in Log Delivery Service files with the number of requests recorded by EdgePrism.

Page Layout

The report displays data on a single page in these sections:

Data Latency

Data Completeness

Data Transfer

Report Specifications

<table>
<tr><td>Latency</td><td>Latency depends on the selected granularity.<table><tr><th>Granularity</th><th>Latency</th></tr><tr><td>10 Minutes</td><td>5 to 10 minutes</td></tr><tr><td>Hour</td><td>1 hour + delta[1]</td></tr><tr><td>Day</td><td>1 day + delta[1]</td></tr></table>[1]delta = approximately 5-10 minutes</td></tr>
<tr><td>Granularity</td><td>Data Latency Section:
Up to 30 minutes
30-40 minutes
40-50 minutes
More than 60 minutes.</td></tr>
</table>

	Data Completeness and Data Transfer Sections: • 10 min • hour • day
Dimensions	Accounts, Storage Locations
Metrics	Data Latency, Data Completeness, Data Transfer
Delivery Mechanism	EdgeQuery
Associated API Endpoint (s)	• `POST /reporting-api/lds` returns LDS report data according to the filters passed • `POST /reporting-api/lds/data-completeness` returns LDS completeness data according to the filters passed

Selecting a Date Range and Time Zone

Select a date range and timezone in the drop-down menus at the upper right part of the page.

The selected date range influences the data in each chart on the page.

Date Range	Increment
Three days or less	5-minute data increments
Greater than three days and less than or equal to one month	1-hour data increments
Larger than one month	1-day data increments.

Selecting an Account

Select one or more accounts in the **LDS Overview for** drop-down menu at the top of the page.

The selected account influences the data in each chart on the page.

Notes:
- If you don't select an account, the accounts for the default company are automatically selected, and a warning is displayed.
- If you modify the selections in the **LDS Overview for** drop-down menu, the filter selections you made are reset to defaults, and the chart updates automatically to match the defaults.

Selecting Storage Locations

Make one or more selections in the drop-down menu to the right of the accounts drop-down menu. Data in the charts on the page is restricted to the selections you make.

Data Latency Section

This section displays the percentage of request entries delivered in LDS files over four latency (see see Concepts) windows: **up to 30min**, **30-40min**, **45-60min**, **more than 60 min**.

For example, you can view the percentage of requests logged within 30 minutes of the time the requests were received.

You can also view latency measurements by moving your mouse pointer over the chart. Information displays in popups.

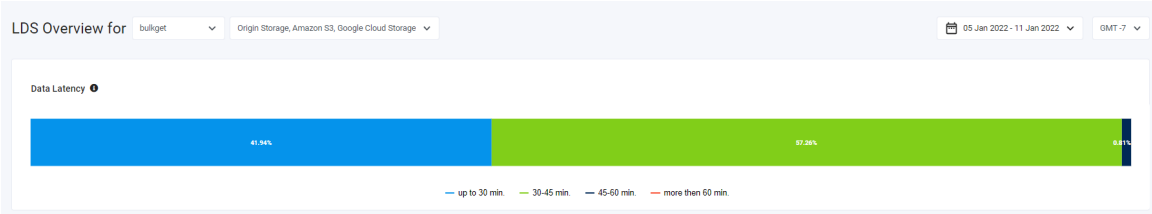

The sum of all percentages is 100%.

Labels beneath the chart allow you to select specific latency time windows.

Toggling Chart Data

The labels are toggles that you can use to display or hide the corresponding chart data. By default, all labels are toggled on. Labels that are toggled off have a gray font color.

Click a label to hide or show the corresponding data in the chart. The chart content is updated to reflect your selection.

Data Completeness Section

This section shows data completeness measurements (see Concepts) in a bar chart.

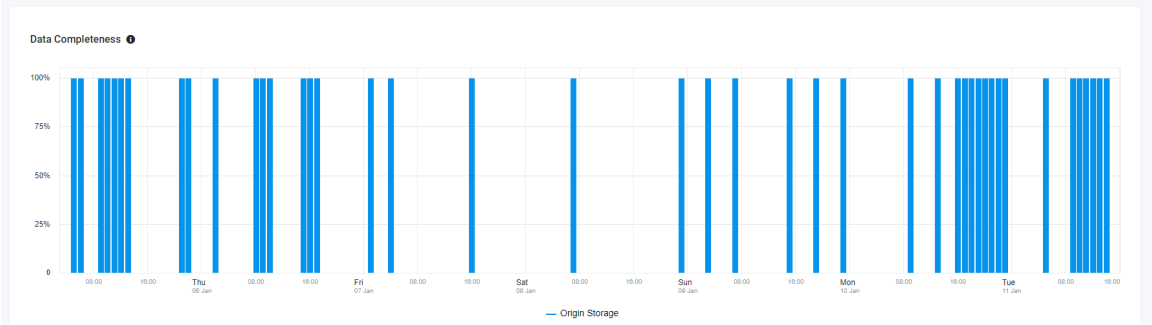

You can also view data completeness measurements by moving your mouse pointer over the chart. Information displays in popups.

Following is the data on the chart:

- X-Axis: Request time in 10-minute, hourly, or daily granularity depending on the date range selected in the drop-down menu at the top right of the report.
- Y-Axis: Percentage of data already uploaded to the storage destination.

Columns are displayed for every destination (Origin Storage, Amazon S3, and so on) for the accounts you selected as described in Selecting an Account. Each destination is identified by a label under the chart.

Note: Data completeness measurements can cross time boundaries. For example due to network issues LDS might deliver only 40 out of 200 requests that occurred during some day for a customer, so the completeness will be 20%. Once the issue is fixed and LDS delivers the rest of the data for that day, completeness will be increased to 100%.

Toggling Chart Data

A legend below the chart identifies the chart content by color. The labels reflect the locations configured for Log Delivery Service configurations.

The legend labels are toggles that you can use to display or hide the corresponding chart data. By default, all labels are toggled on. Labels that are toggled off have a gray font color.

Click a label to hide or show the corresponding data in the chart. The chart content is updated to reflect your selection.

Data Transfer Section

This section shows the volume of data transferred to your storage location.

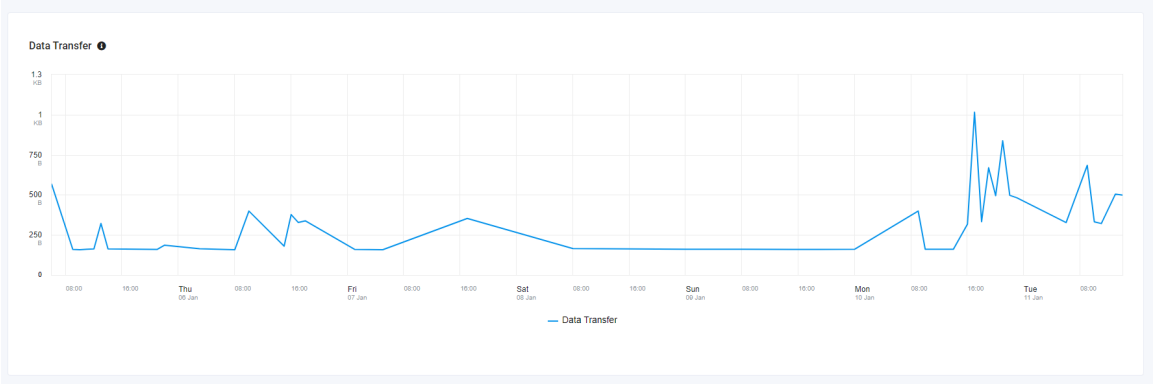

Following is the data on the chart:

- X-Axis:Delivery time in 10-minute, hourly, or daily granularity depending on the date range selected in the drop-down menu at the top right of the report.
- Y-Axis: Size of delivered files in gigabytes (GB).

Toggling Chart Data

A legend below the chart identifies the chart content by color and has a single entry, **Data Transfer**.

The legend labels are toggles that you can use to display or hide the corresponding chart data. By default, all labels are toggled on. Labels that are toggled off have a gray font color.

Click a label to hide or show the corresponding data in the chart. The chart content is updated to reflect your selection.

How Metrics Are Calculated

Metric	Calculation
Data Latency	Request Delivery Time - Request Logging Time Example: A request was logged by EdgePrism at 10:30 and Log

Metric	Calculation
	Delivery Service delivered it in file at 10:55, so delivery latency for that request is 25 minutes.
Data Completeness	Number of log lines delivered by LDS / number of egress requests from the reporting system. The maximum value is 100%.
Data Transfer	No calculation. Data provided by EdgeQuery.

Live Push Report

The Live Push report provides many facets of information on Live Push Ingest from your publisher application to the LimelightCDN.

You can use the report to:

- Ascertain the availability of your live content to be delivered over the CDN
- Assist in determining the stability and efficiency of your publisher application and "first-mile" transit via request time and response codes
- View the storage (number of files, total size) used by your chunked streaming files
- Monitor aspects of your stream ingests by account, slot, and more

> **Note:** "First-mile" is the connection from your publisher application to the Limelight CDN.

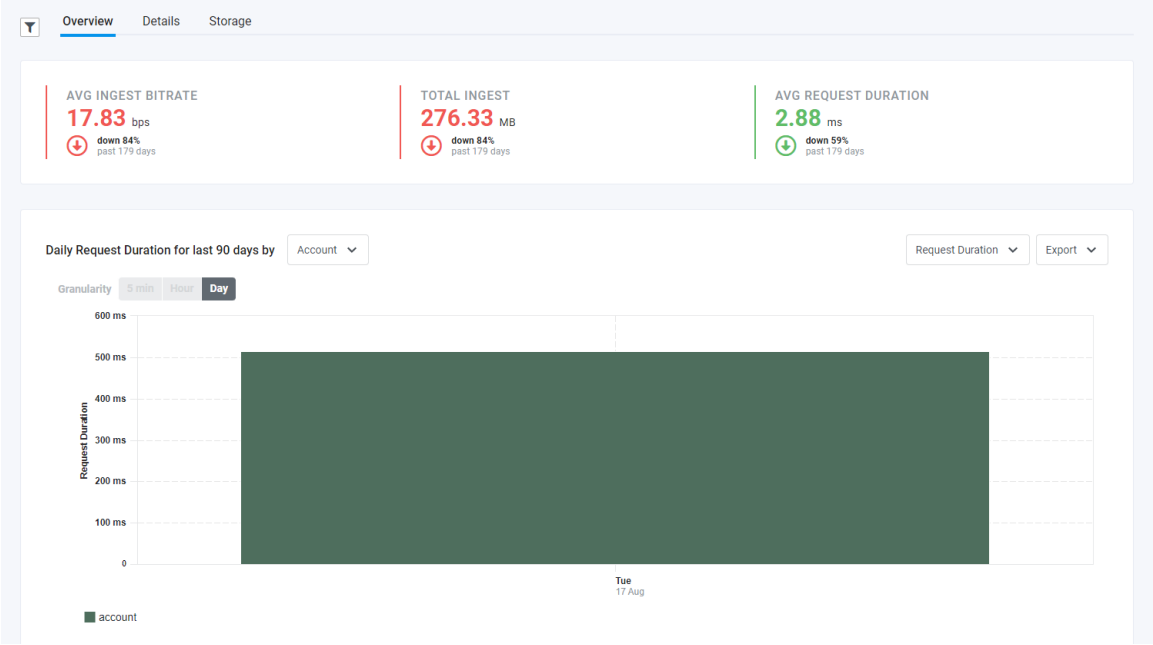

The report has these tabs:

Overview

Details

Storage

Report Specifications

Latency	One hour
Granularity	5 minutes, hour, day **Note:** For day granularity, data is collected in the GMT-7 timezone.
Dimensions	Account, Stream Name, Status Code

Metrics - Summary Area	Average Ingest Bitrate, Total Ingest, Average Request Duration
Metrics - Chart	Data Transfer In, Data Transfer Out, Request Duration
Metrics - Details Tab	Ingest Bytes, Egress Bytes, Requests, Request Duration
Metrics - Storage Tab	Total Bytes Stored, Total Files
Delivery Mechanism	Billing API
Associated API Endpoint(s)	• `GET /billing/v2/shortnames/{short name}/LivePushIngest` Returns Live Push Ingest billing information • `PUT https://{shortname}-{slotname}-pri.live-push.llnw.net` Puts a chunked streaming file onto the Live Push Ingest server

Selecting Accounts, Date Range, and Time Zones

You can make selections in the controls above the tab header:

- **LIVE PUSH FOR**. Select one or more accounts to which you have access for cross-account analysis. Click the **Select All** button to select all accounts.
- **Date range**. Pick from pre-set time frames or choose custom date ranges in the drop-down menu. Click the **Apply** button on custom ranges.
- **Time zone**. The top five most commonly used timezones in Control are at the top of the drop-down menu. Scroll down for additional time zones.

Choosing Stream Names and Status Codes

Customer publishing applications send requests to store chunked streaming files on the Live Push Ingest server, and the server returns status codes to indicate the state of the request.

Use the **Stream Name/Status Codes Filter** to determine which stream (slot) names and status codes to display in all the report's tabs.

The filter icon (a funnel) on the tab header's left side allows you to filter by stream names and status codes. The icon is a toggle that displays or hides the filter control.

1. Click the filter icon.
2. Make selections in the subsequent **Filter by**drop-down menu.
   By default, the **Stream Name** and **Status Code** entries are unchecked. As such, data for all streams and status codes are displayed.
3. If desired, select **Stream Name** or **Status Code** to display additional drop-down menus to further filter for display in the chart.
4. Click the **Apply** button.

Overview Tab

This tab presents a synopsis of Live Push data for the selected account and reporting data range. Included are a summary area and a stacked bar chart.

Summary Area

For various statistics, the Summary Area shows the percent change for the selected reporting date range relative to the prior time period of the same length. For example, if you select **This Month**, the statistics show a comparison between this month's data and the previous month's.

Colors and arrows represent changes:

- An increase displays in green with an arrow pointing up.
- A decrease displays in red with an arrow pointing down.
- If there was no change from the previous period, the text is gray with an empty circle instead of an arrow.
- Percentage up or down is in bold text to the right of the arrow or empty circle.
- Information in gray text beneath the percentage up or down indicates the total date range covered minus any time remaining in the current period. For example, if you select **Last 7 days** as the reporting period and the current date is part of the past seven days, the past 13 days' information is displayed.

Statistics in the *Summary Area* are:

Statistic	Description
Average Ingest Bitrate	Average ingest bitrate of all data coming into the Live Push Ingest server from your publisher application, measured in bits per second.
Total Ingest	Total of all data coming into Live Push Ingest server from your publisher application, measured in bytes.
Average Request Duration	Average duration of requests from your publisher application to store a chunked streaming file on a Live Push Ingest server, where a request duration encompasses the time request arrives at the server to the time the request is complete. Measured in seconds.

Selecting a Data Grouping

Beneath the Summary Area on the left is a drop-down menu that allows you to determine how data is broken out in the chart: by Account, Stream Name, or Status Code.

Make a selection.

Selection	Data Displayed in Chart	Data Displayed in Legend
Not Selected	All data for the selected metric.	Single entry: **All Data**.
Account	Data for the selected account and metric.	A label for each account chosen.
Stream Name, Status Code	Data for the filtered Stream Name and Status Code.	A label for each Stream Name and Status Code you filtered.

The chart reflects your selection.

See also Choosing Metrics, Choosing Accounts, and Choosing Stream Name and Status Codes.

Choosing Cart Granularity

The chart can be further refined by selecting one of the **Increment** values:

- 5 min
- Hour
- Day

The selection you make determines increments along with Y-axis. For report date ranges of 3 days or less, 5-minute data increments are displayed. For date ranges greater than 3 days but less than or equal to 1 month, 1-hour data increments are shown. For larger date ranges, 1-day data increments are displayed.

Choosing Metrics

Above the chart, on the right, is a drop-down menu that allows you to select chart metrics.

1. Make a selection.

Selection	Description
Ingress Bytes	Total of all data coming into Live Push Ingest server from your publisher application, measured in bytes.
Egress Bytes	Total data of all your streams egressing from the Live Push Ingest server to the CDN.
Request Duration	Time taken to process a request to store a chunked streaming file on a Live Push Ingest server, from the time a request arrives at the server to the time the request is complete.

2. The chart reflects your selection.

Toggling Chart Data

A legend beneath the chart identifies the chart content by color. The legend reflects the choices you made in the **Stream Name/Status Code Filter** and the **Data Breakout**drop-down menu.

For example, if you filtered by two specific stream names and selected **Stream Name** in the drop-down menu, the stream's data is displayed in the chart. Labels for the two names display in the chart legend.

The legend labels are toggles that you can use to display or hide the corresponding chart data. By default, all labels are toggled on. Click a label to hide or show the corresponding data in the chart. Labels that are toggled off have a gray font color.

Exporting Chart Data

To export data currently displayed in the chart, click the **Export** drop-down menu on the right above the chart; then choose an option:

- **PowerPoint**: Export to a PowerPoint file that contains a screenshot of the chart.
- **CSV**: Export data to a Comma-Separated Values (CSV) file that reflects the currently selected chart options.

After you choose an option, Control creates and downloads the report.

Details Tab

This tab shows Live Push Ingest metrics in a tabular format. Rows in the table can be expanded or collapsed, allowing you to drill down into the table. You can drill down in the table to view metrics by month, day, and hour. The metrics displayed are:

Metric	Description
Ingest Bytes	Total of all data coming into Live Push Ingest server from your publisher application, measured in bytes.
Egress Bytes	Total data of all your streams egressing from the Live Push Ingest server to the CDN.
Requests	Total number of requests from your publisher application to the Live Push Ingest server.
Request Duration	Duration of requests from your publisher application to the Live Push Ingest server to store chunked streaming files on the Live Push Ingest server.

Selecting a Data Grouping

Above the table on the left is a drop-down menu you can use to set the table's grouping. Make a selection.

Selection	Data Grouped by
Not Selected	All data for the selected metric.
Account	The account you selected above the tab header.
Stream Name	The stream names associated with the account.
Status Code	Status code: 201, 204, 400.

The first-level groupings reflect the item you selected. For example, if you select **Status Code**, the table displays three expandable rows, one for each status code.

If you chose **Not Selected**, data is grouped by month under a single expandable heading: **All Data**.

Click a first-level grouping to display metrics for months. Click a month to display metrics for the days in the month. Click once more to display metrics for each hour in the day.

Notes:

- Rows that include "(partial)" indicate not all data is available for the time period. Only status codes returned during the selected time frame are displayed.
- Only status codes returned for the time period are displayed.

Exporting Data

You can export data in the table to a Comma-Separated Values (CSV) file.

Click the **Export CSV**drop-down menu on the right above the table, then choose an option.

Option	CSV File Contains
Monthly	A total for each metric in each month.
Daily	A total for each metric in each day in each month.
Hourly	A total for each metric in each hour in each day in the month.

After you choose an option, Control creates and downloads the report.

Storage Tab

This tab shows metrics about your chunked streaming files stored on the Live Push Ingest server.

Selecting a Data Grouping

Above the chart on the left is a drop-down menu you can use to set the grouping the in table. Make a selection.

Selection	Data Grouped by
Not Selected	All data for the selected metric.
Account	The account you selected above the tab header.
Stream Name	The stream names associated with the account.

Choosing Metrics

On the right, above the chart, is a drop-down menu that allows you to select metrics about your chunked streaming files stored on the Live Push Ingest server.

Make a selection.

Selection	Description
Total Bytes Stored	Total size of all files.
Total Files	Total number of files.

How Metrics Are Calculated

The Overview Tab Summary Area and chart show values for the chosen metric in the **Stream Name/Status Codes Filter**.

The unit of measurement depends on the chosen metric.

Overview Tab - Summary Area

Metric	Calculation
Average Ingest Bitrate	The average bitrate of the stream used to ingest content.
Total Ingest	No calculation. Data provided by EdgeQuery.

Metric	Calculation
Average Request Duration	The average duration of requests.

The up or down change for a given metric is calculated as follows.

The metric from the previous period is compared to the metric from the current period.

Comparison Results	Metric Presentation
The selected time period's value equals the previous time period.	Presented as no change.
The previous period's value is smaller than the current period.	Presented as up from the previous period. See also Percentage Calculation.
The previous period's value is greater than the current period.	Presented as down from the previous period. See also Percentage Calculation.

Percentage Calculation

The 'Up' and 'Down' presentations include a percentage.

Percentage = ((newValue - oldValue) / oldValue) * HUNDRED_PERCENT

Overview Tab - Chart

Metric	Calculation
Data Transfer In	No calculation. Data provided by EdgeQuery.
Data Transfer Out	No calculation. Data provided by EdgeQuery.
Request Duration	Duration of the request.

Details Tab

Metric	Calculation
Ingest Bytes	No calculation. Data provided by EdgeQuery.
Egress Bytes	No calculation. Data provided by EdgeQuery.
Requests	No calculation. Data provided by EdgeQuery.
Request Duration	No calculation. Data provided by EdgeQuery.

Live Stats Report

The Live Stats Report shows real-time bandwidth usage and can be used to monitor real-time traffic trends.

The report displays usage in *Bits/Sec* for one or more services for the preceding 6 hours.

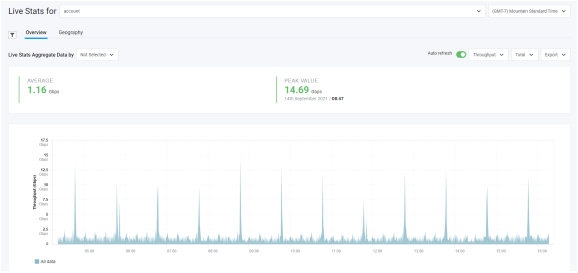

The report automatically refreshes data every 30 seconds but contains an **Auto refresh** toggle in both tabs that allows you to enable or disable automatic refresh. The selection you make in one tab is reflected in the other tab.

The report has these tabs:

Overview

Geography

Report Specifications

Latency	One minute
Granularity	One minute
Dimensions	Account, Protocol,
Metrics - Summary Area in Overview Tab	Average, Peak Value
Metrics - Chart in Overview Tab, **Map in Geography Tag**	Throughput In and Out Data Transfer In and Out Requests In and Out
Delivery Mechanism	Realtime Reporting API
Associated API Endpoint(s)	• `/realtime-reporting-api/traffic/livestats` Retrieves report data based on the filters applied. • `/realtime-reporting-api/traffic/livestats/geo` Retrieves geo report data based on the filters applied.

Choosing Accounts

You can choose accounts for which you wish to see live stats information. Click the drop-down menu to the right of the *LIVE STATS FOR* title and select the desired accounts.

Note: You must select at least one name; otherwise, the default company is automatically selected, and a warning is displayed.

Choosing a Time Zone

Use the **Time Zone Control** at the top right of the page to select a timezone. The top five most commonly used timezones in Control are at the top of the drop-down menu. Scroll down for additional time zones.

Overview Tab

This tab summarizes your data in a summary area and a chart with a legend to its right.

Filtering for Protocols to Display

Use the **Protocol Filter** to determine which protocols to display in the chart.

The filter icon (a funnel) on the left side of the tab header allows you to filter by protocol. The icon is a toggle that displays or hides the filter control.

1. Click the filter icon.
2. Make a selection in the subsequent **Filter by**drop-down menu.
   By default, the **Protocol** entry is unchecked. As such, data for all protocols is displayed in the chart.
3. If desired, select **Protocol** to display a drop-down menu with protocols you can select for display in the chart.
4. Click the **Apply** button.

Selecting a Data Grouping

On the left under the tab header is the **Live Stats Aggregate Data by**drop-down menu that allows you to determine how data is broken out in the chart: by Account or Protocol.

Make a selection:

Selection	Data Displayed in Chart	Information Displayed in Legend
Not Selected	All data for the selected metric.	Single entry: **All Data**.
Account	Data for the accounts you selected (see Choosing Accounts).	List of all accounts you selected.
Protocol	Data for all protocols associated with the accounts you selected.	List of protocols.

The chart reflects your selection, and the Summary Area displays data for the selected metric.

Choosing Metrics

See Choosing Metrics for the Overview and Geography Tab.

Summary Area

The *Summary Area* summarizes the metrics you selected (see Choosing Metrics for the Overview and Geography Tab) for the last six hours.

Units depend on the selected direction (In, Out, Total) and metric:

- Throughput - bits per second
- Data Transfer - bytes
- Requests - number of requests

Statistic	Description
Average	Average for the selected metric and direction.
Peak Value	Highest value of the selected metric and direction.

Exporting Chart Data

To export data currently displayed in the chart, click the **Export** drop-down menu on the right above the chart; then choose an option:

- **PowerPoint**: Export to a PowerPoint file that contains a screenshot of the chart.
- **CSV**: Export data to a Comma-Separated Values (CSV) file that reflects the currently selected chart options.

After you choose an option, Control creates and downloads the report.

Toggling Chart Data

A legend beneath the chart identifies the chart content by color. The legend reflects the choice that you made in the accounts drop-down menu (Choosing Accounts) and the aggregate drop-down menu (Selecting a Data Grouping). For example, if you selected 'Protocol' in the aggregate drop-down menu, the legend contains a list of all protocols associated with the account.

The legend labels are toggles that you can use to display or hide the corresponding chart data. By default, all labels are toggled on. Click a label to hide or show the corresponding data in the chart. Labels that are toggled off have a gray font color.

Geography Tab

This tab shows usage in *Bits/Sec* over the last 12 hours on an interactive map.

> **Note:** The map shows data only if your account has the 'POP view' and 'Region view' products.

In the *Region Data* chart, individual POP traffic is stacked for each region, and you can mouse over any bar segment to view POP traffic details.

Click a continent to view data at the country level.

Choosing Metrics

See Choosing Metrics for the Overview and Geography Tab.

Exporting Data

You can export data to a Comma-Separated Values (CSV) file that reflects the currently selected map metrics.

Click the **Export**drop-down menu on the right above the map, then select **CSV**.

After you select **CSV**, Control creates and downloads the report.

Choosing Metrics for the Overview and Geography Tab

On the right under the tab header are drop-down menus that allow you to select metrics (Throughput, Data Transfer, Requests) and direction (In, Out, Total).

Make selections:

Metric	Metric Description	Direction
Throughput	Traffic flow measured in bps.	In: Through the network from the origin. Out: through the network to the requestor,
Data Transfer	Data flow measured in bytes.	In: Into the CDN from the origin. Out: From the published host to the requestor.
Requests	Number of requests.	In: Into the CDN from the origin. Out: From the published host to the requestor.

Note: For all metrics, Total = In + Out.

The metrics you select are reflected in the summary area and chart in the Overview tab, and the map in the Geography tab.

How Metrics Are Calculated

Overview Tab - Summary Area

Metric	Calculation
Average Throughput	Average of all throughput values for the reporting period.
Average Data Transfer	Average of all data transfer values for the reporting period.
Average Requests	Average of all request values for the reporting period.
Peak Value Throughput, Data Transfer, Requests	Max value of the selected metric for the reporting period.

Overview Tab Chart and Geography Tab Map

Metric	Calculation
Throughput In	No calculation. Data provided by EdgeQuery.
Throughput Out	No calculation. Data provided by EdgeQuery.
Data Transfer In	No calculation. Data provided by EdgeQuery.

Metric	Calculation
Data Transfer Out	No calculation. Data provided by EdgeQuery.
Requests In	No calculation. Data provided by EdgeQuery.
Requests Out	No calculation. Data provided by EdgeQuery.

Realtime Streaming Report

The *Realtime Streaming Report* allows you to monitor the performance and popularity of your Realtime Streaming streams. The report provides a history of client edge connection data.

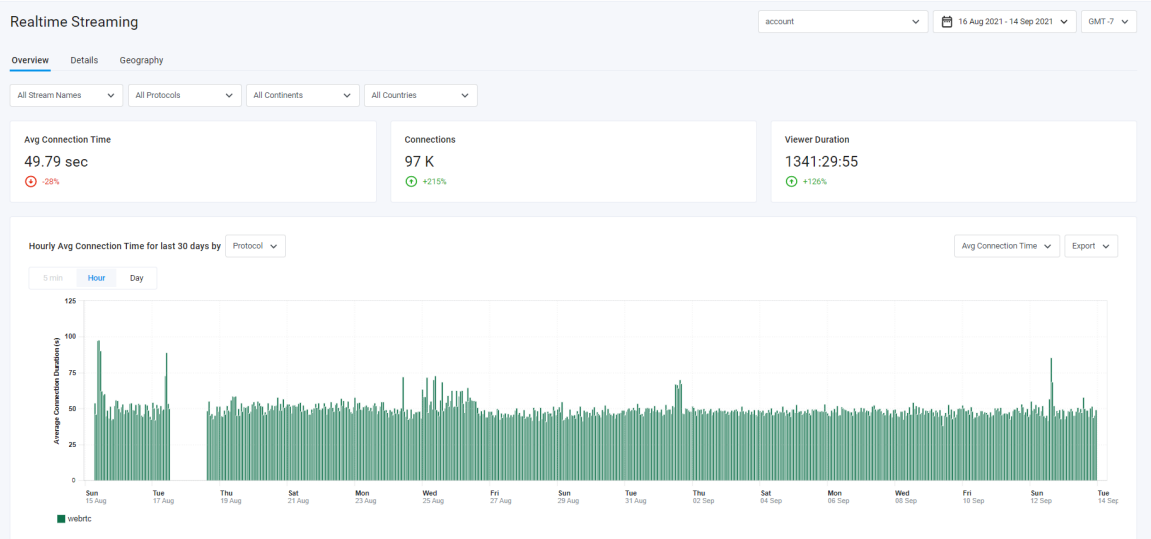

The report has these tabs:

Overview

Details

Geography

Report Specifications

Latency	Latency depends on the selected granularity. **Granularity** / **Latency** 5 Minutes / 5 to 10 minutes Hour / 1 hour + delta* Day / 1 day + delta* *delta = approximately 5-10 minutes
Granularity	5 minutes, hour, day
Dimensions	Slot (stream name), Protocol, Continent, Country
Metrics	Average Connection Time, Number of Connections, Viewer Duration
Delivery Mechanism	RealTime Streaming Subscriber Logs
Associated	Examples

API Endpoint (s)	• `/realtime-reporting-api/edgefunctions/` returns EdgeFunctions traffic data • `/realtime-reporting-api/edgefunctions/geo` returns EdgeFunctions traffic geographic data

Selecting Accounts, Date Range, and Time Zones

You can make selections in the controls above the tab header:

- **REALTIME STREAMING FOR**. You can choose accounts for which you wish to see Realtime Streaming data. Click the drop-down menu and select the desired accounts. The chart data is limited to the accounts you select.
- **Date range**. Pick from pre-set time frames or choose custom date ranges in the drop-down menu. Click the **Apply** button on custom ranges.
- **Time zone**. The top five most commonly used timezones in Control are at the top of the drop-down menu. Scroll down for additional time zones.

Selecting Report Dimensions

Beneath the tab headers are drop-down menus that allow you to select the dimensions to display in the chart on the **Overview** tab and the table in the **Details** tab.

Dimension	Description
Stream Name*	Stream names configured for each selected account.
Protocol*	Realtime Streaming, RTMP, and so on.
Continents, Countries	Geographical regions.

*Not available in the Geography tab.

1. Select the desired entries in each drop-down menu, or leave unselected to include all entries in a drop-down menu.

   Note: **Continents** and **Countries** are mutually exclusive.

2. Click the **Apply** button.

Overview Tab

This tab presents a synopsis of your Realtime Streaming streams for the selected accounts and reporting data range. Data is presented in the *Summary Area* and in a stacked bar chart.

Summary Area

For various statistics, the Summary Area shows the percent change for the selected reporting date range relative to the prior time period of the same length. For example, if you select **This Month**, the statistics show a comparison between this month's data and the previous month's.

Colors and arrows represent changes:

- An increase displays in green with an arrow pointing up.
- A decrease displays in red with an arrow pointing down.
- If there was no change from the previous period, the text is gray with an empty circle instead of an arrow.
- Percentage up or down is in bold text to the right of the arrow or empty circle.

- Information in gray text beneath the percentage up or down indicates the total date range covered minus any time remaining in the current period. For example, if you select **Last 7 days** as the reporting period and the current date is part of the past seven days, the past 13 days' information is displayed.

> **Note:** Information in the *Summary Area* depends on the selected accounts, time range, time zone, and dimensions.

Statistics in the *Summary Area*

Statistic	Description
Average Connection Time	The average session duration of client connections open across streams during the time period. Measured in seconds.
(Number of) Connections	Total number of connections from clients to your Realtime Streaming streams during the time period.
Viewer Duration	Count of how many viewer hours were logged during the time period. Measured in hours, minutes, and seconds.

Selecting a Data Grouping

Use the drop-down menu on the left beneath the *Summary Area* to select a dimension to group data. The selected dimension is reflected in the chart content.

Selecting Metrics

On the right beneath the *Summary Area* is a drop-down menu that allows you to select chart metrics. The selected metric is reflected in the y-axis label. See Selecting Metrics (All Tabs) for information.

After you make a selection, the chart refreshes to reflect your selection.

Choosing Granularity

The chart can be further refined by clicking one of the **Increment** toggle buttons on the left above the chart. Buttons are enabled depending on the date range you selected.

- **5 min**: for ranges 24 hours or less.
- **Hour**: for ranges less than 14 days.
- **Day**: for all ranges.

Toggling Chart Data

A legend below the chart identifies the chart content by color. The labels reflect the dimensions in the chart.

The legend labels are toggles that you can use to display or hide the corresponding chart data. By default, all labels are toggled on. Labels that are toggled off have a gray font color.

Click a label to hide or show the corresponding data in the chart. The chart content is updated to reflect your selection.

Exporting Chart Data

To export data currently displayed in the chart, click the **Export** drop-down menu on the right above the chart; then choose an option:

- **PowerPoint**: Export to a PowerPoint file that contains a screenshot of the chart.
- **CSV**: Export data to a Comma-Separated Values (CSV) file that reflects the currently selected chart options.

After you choose an option, Control creates and downloads the report.

Details Tab

This tab shows metrics in tabular form by month, and metrics are displayed in a table with expandable rows.

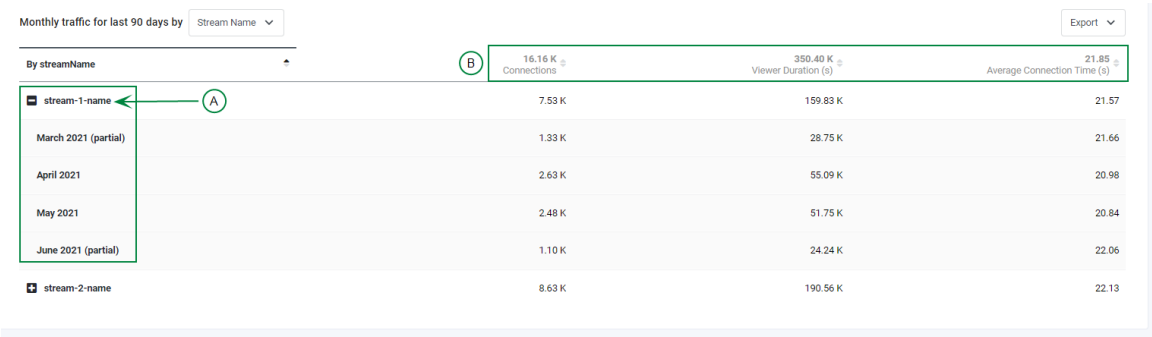

In the preceding figure:

A - Expanded row with data broken out by month

B - Metrics

Selecting a Data Grouping

Use the drop-down menu on top left of the table to select a dimension to group data in the table Information in the table is updated to reflect your selection.

Working with Rows in the Table

Metrics are shown in tabular format as expandable rows where the expandable item depends on the data grouping you selected. For example, if you had selected **Protocol**, then the expandable items would be **rtmp**, **webrtc**, and so on.

Click on an expandable item to reveal metrics by month. The range of months depends on the date range you selected. See Selecting Accounts, Date Range, and Time Zones. For a description of metrics, see Selecting Metrics (All Tabs).

Removing and Adding Columns

1. Hover the mouse pointer over a column header to reveal an **x** icon.
2. Click the icon to remove the column.
   The column is removed from the table.

When additional columns are available to display, a **+** icon displays on the right side of the table header row.

1. Click the + icon to reveal a menu with additional columns.
2. Click a column to add.
   The column is added to the table.

Exporting Table Data

To export chart data to a Comma-Separated Values (CSV), click the **Export** drop-down menu on the right above the table; then select **CSV**.

After you choose **CSV**, Control creates and downloads the report.

Geography Tab

This tab shows metrics on an interactive map.

Selecting a Data Grouping

Use the drop-down menu on top left of the map to select a dimension to group data. The selected dimension is reflected in the map content.

Selecting Metrics

On the right beneath the *Summary Area* is a drop-down menu that allows you to select metrics. See Selecting Metrics (All Tabs) for information.

Viewing Data

- Hover the mouse pointer over a content to view continent-level data.
- Click a continent then hover over countries in the continent to view data at the country level.
- Click the **Whole World** button (under the metric drop-down menu) to zoom out to world level.

Selecting Metrics (All Tabs)

In all tabs, drop-down menus on the right beneath the *Summary Area* allow you to choose metrics. Available metrics depend on the tab; all metrics are described here.

Metric	Description
Viewer Duration	Count of how many viewer hours were logged during the time period. Measured in hours, minutes, and seconds.
Connections	Total number of connections from clients to your Realtime Streaming streams during the time period.
Average Connection Time	The average session duration of client connections open across streams during the time period.

How Metrics Are Calculated

Metrics are calculated as described in the following table.

Metric	Calculation
Average Connection Time	The average session duration of client connections open across

Metric	Calculation
	streams during the time period.
Connections	No calculation. Data provided by EdgeQuery.
Viewer Duration	Sum of viewer hours logged during the time period.

Overview Tab - Summary Area

The up or down change for a given metric (Average Connection Time, Connections, or Viewer Duration) is calculated as follows.

The metric from the previous period is compared to the metric from the current period.

Comparison Results	Metric Presentation
Selected time period value equals previous time period.	Presented as no change.
Previous period value is smaller than current period.	Presented as up from previous period. See also Percentage Calculation.
Previous period value is greater than current period.	Presented as down from previous period. See also Percentage Calculation.

Percentage Calculation

The 'Up' and 'Down' presentations include a percentage.

Percentage = ((newValue - oldValue) / oldValue) * HUNDRED_PERCENT

Geography Tab

The chosen metric is calculated for the chosen continent or country:

- Country level: sum of the metric in the country.
- Continent level: sum of the metric for all countries in the continent.

Service Provider Traffic Report

The Service Provider Traffic Report provides metrics and trends that allow service providers to:

- Identify, isolate, and resolve performance-related issues to maintain a superior end-user experience and potential revenue share from Limelight.
- Proactively identify potential issues and take corrective actions or escalate to Limelight Operations for further troubleshooting from a software perspective.

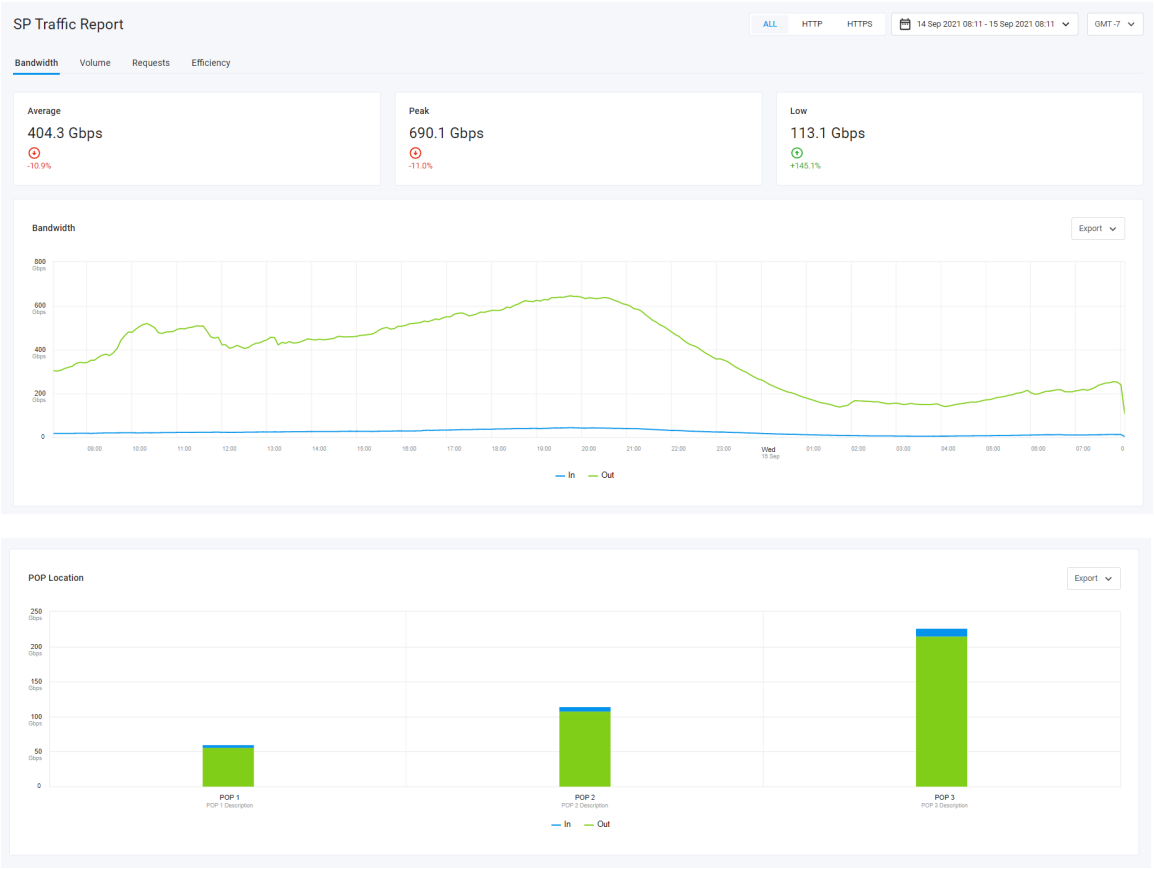

Figure 1. Service Provider Traffic Report

Report Specifications

Latency	Five minutes
Gran-ularity	**Selected Time Frame** / **Granularity** Today Last 24 (days) Last 7 (days) / 5 minutes Last month Last 30 (days) / hourly

	<table><tr><th>Selected Time Frame</th><th>Granularity</th></tr><tr><td>Last 90 (days) Any custom date range > 90 days</td><td>daily</td></tr></table>
Dimensions	• Date/time • In, Out • Limelight content provider aggregate • POP • Service provider account • Service (HTTP, HTTPS) • Service provider's content providers
Metrics	Average, Peak, and Low values for Bandwidth, Volume (data transferred), and Requests
Delivery Mechanism	EdgeQuery

Choosing Services, Time Frames, and Timezone

The top right part of the report contains controls for selecting services, a time frame, and a timezone.

- Choose a service: HTTP, HTTPS, or All (HTTP and HTTPS).
- Choose from pre-set time frames or choose custom date ranges in the drop-down menu. Click the **Apply** button on custom ranges.
- Choose a timezone. The top five most commonly used timezones in Control are at the top of the drop-down menu. Scroll down for additional time zones.

The data in the tabs changes to reflect your choices.

Metrics in the Report

The tabs in the report provide the metrics shown in the following table:

Tab	Metric
Bandwidth	Data rate, measured in bytes per second
Volume	Amount of data transferred, measured in bytes
Requests	Number of requests for data
Efficiency	Volume of data and the number of requests served from the cached versus the volume of data and number of requests that resulted in a cache miss at the PoP level.

Tab Components

The following sections describe the **Summary Area** and various charts available in the Service Provider Traffic report.

Summary Area

> **Note:** Only the **Bandwidth**, **Volume**, and **Requests** tabs contain a Summary Area.

This component shows the average value, peak value, and low value for the metric in the selected tab, along with a comparison to the prior period of the same time frame.

- The value is displayed in bold black text with the appropriate unit of measurement; for example, bytes per second for bandwidth.
- A percentage up or down from the previous period is shown in green, red, or gray:
  - Increase: a circle with an arrow pointing up along with the percentage change, both in green.
  - Decrease: a circle with an arrow pointing down, along with the percentage change, both in red.
  - No change: a circle with an arrow pointing to the right, 0% as the percentage change, circle and percentage both in blue.

The percentage up or down is calculated as follows:

Percentage = ((newValue - oldValue) / oldValue)) * HUNDRED_PERCENT

Chart Area

The chart area provides a graphical representation of the measurement represented in a tab.

Line Chart

POP Location Chart

Volume Distribution by Content Provider Chart

POP Location Data Transfer Efficiency Chart

Line Chart

> **Note:** Only the **Bandwidth**, **Volume**, and **Requests** tabs contain the line chart.

The line chart shows In and Out values for the measurement represented in the tab.

- In: any traffic coming into the PoP from other PoPs or customer origins
- Out: any traffic leaving the PoP from the edge to requesting clients

The In and Out values are a function of the selected time frame. For example, a time frame of 24 hours shows hours on the X-axis.

Move your pointer across the chart to view details over time.

POP Location Chart

> **Note:** Only the **Bandwidth**, **Volume**, and **Requests** tabs contain the PoP Location bar chart.

This stacked bar chart summarizes all In and Out values, broken out by PoPs, where each bar represents a PoP location.

- To view In values and Out values, hover your pointer over a bar.
- The PoP name and description are beneath each bar.

Volume Distribution by CPs Chart

> **Note:** Only the **Volume** tab contains this chart.

This chart summarizes content provider traffic volume for each of the service provider's content providers. Each bar represents one content provider, and the 'In' and 'Out' values are aggregates of all the content provider's accounts (shortnames)

- Scroll to the bottom of the page to view the chart.
- To view 'In' and 'Out' values, hover your pointer over a bar.

POP Location Data Transfer/Requests Efficiency Chart

> **Note:** Only the **Efficiency** tab contains this chart.

This chart shows the efficiency of a PoP in terms of volume of data and number of requests served from the cache versus those that resulted in a cache miss at the PoP level.

Data is shown in percentages where higher values indicate higher efficiency.

- Use the **Data Transfer | Requests** toggle above the chart to display information for Data Transfer or Requests.
- To view the efficiency measurement, hover your pointer over a bar.
- The PoP name and description are beneath each bar.

Data Transferred Efficiency Calculation	Requests Efficiency Calculation
(bytes served from cache - bytes served from origin) / (bytes served from cache) * 100%	(requests served from cache - requests served from origin) / (requests served from cache) * 100%

> **Note:** If the value of Egress - Ingress is negative, the efficiency value displayed is zero.

Exporting Data

To export chart data to a PowerPoint file that contains a screenshot of the chart, click the **Export**drop-down menu on the right above the chart; then select **PowerPoint**.

After you select **PowerPoint**; Control creates and downloads the report.

Traffic Report

The *Traffic Report* provides an integrated view of traffic for one or more selected Accounts. The report loads with the following smart defaults, which you can change by making further selections in the header.

- The Account is your "default" account, which you set in *My Account > Profile > Default Account*.
- The time frame is "Last 30 days."
- The timezone is your chosen "Timezone" setting from your profile in "My Account."
- The data payload is organized "By protocol" and includes all protocols for which data exists within the selected timeframe and timezone. The report is capable of showing these protocols: DASH, HTTP, HTTPS, MSS, HLS.

Traffic from both IPv4 and IPv6 addresses is reported. If you need IPv6 support but have not previously requested it be enabled for your Limelight Account, please contact Limelight Customer Support.

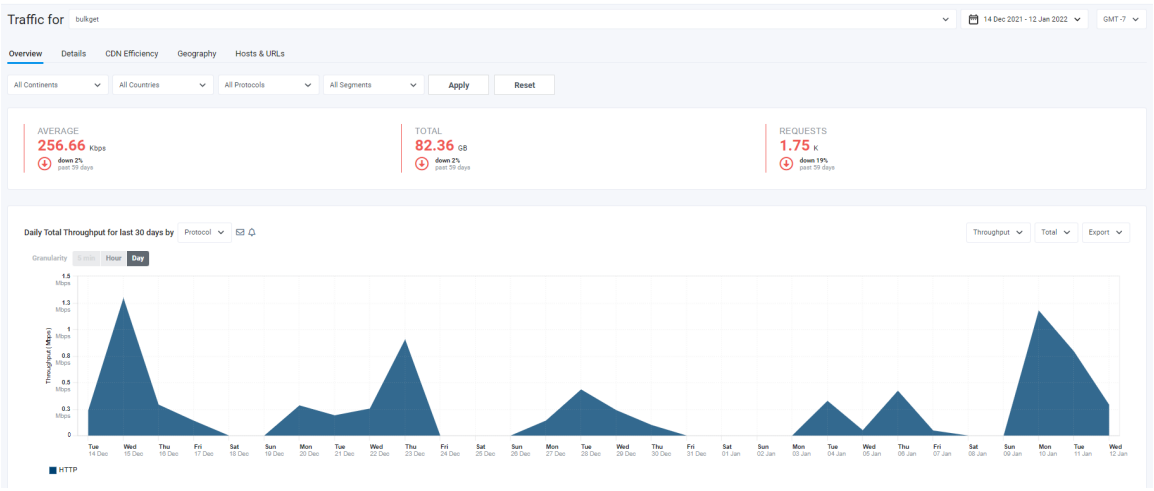

The report has these tabs:

Overview

Details

CDN Efficiency

Geography

Hosts & URLs

Report Specifications

Latency

Latency depends on the selected granularity.

Granularity	Latency
5 Minutes	5 to 10 minutes
Hour	1 hour + delta*
Day	1 day + delta*

*delta = approximately 5-10 minutes

Granularity	5 minutes, hour, day
Dimensions	Account, Protocol, Segment, Continent, Country, CDN Efficiency, URLs, Published Hosts, Referrer URLs, File Types
Metrics - Summary Area	Average traffic rate Total bytes transferred Total requests
Metrics - Elsewhere	Throughput (In, Out, and Total), Data Transfer (In, Out, and Total), or Requests (In, Out, and Total).
Delivery Mechanism	EdgeQuery
Associated API Endpoint (s)	• `GET /realtime-reporting-api/traffic` returns Realtime Reporting API traffic data according to query parameters passed • `POST /realtime-reporting-api/traffic` returns Realtime Reporting API traffic data according to parameters passed in the request body • `GET /realtime-reporting-api/traffic/retentions` returns a list of all retentions applied to traffic reports • `GET /realtime-reporting-api/traffic/continents` returns a list of continents associated with the account • `GET /realtime-reporting-api/traffic/countries` returns a list of countries associated with the account • `GET /realtime-reporting-api/traffic/states` returns a list of countries associated with the account • `GET /realtime-reporting-api/traffic/geo` returns realtime reporting traffic geo data according to query parameters passed • `POST /realtime-reporting-api/traffic/geo` returns realtime reporting traffic geo data according to parameters passed in the request body • `GET /realtime-reporting-api/traffic/retentions` returns a list of all retentions applied to traffic geo reports

Selecting Accounts, Date Range, and Time Zones

You can make selections in the controls above the tab header:

- **TRAFFIC FOR**. Select one or more accounts to which you have access for cross-account analysis. Click the **Select All** button to select all accounts.

  > **Note:** You must select at least one account; otherwise, the default company is automatically selected, and a warning is displayed.

- **Date range**. Pick from pre-set time frames or choose custom date ranges in the drop-down menu. Click the **Apply** button on custom ranges.
- **Time zone**. The top five most commonly used timezones in Control are at the top of the drop-down menu. Scroll

down for additional time zones.

> **Note:** Data in the **Hosts & URLs** tab is of DAY granularity and is always stored in GMT-7, so the **Time zone** drop-down menu is disabled when you select the tab.

Filtering Report Data

You can filter report data using filter controls on the **Overview**, **Details**, **CDN Efficiency**, and **Geography** tabs. You can filter on **Continents**, **Countries**, **Protocols**, and **Segments**.

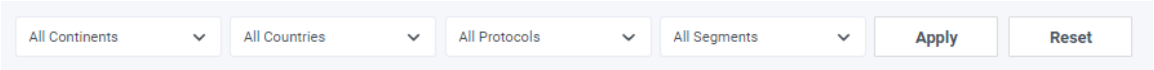

Make the desired selections, then click the **Apply** button to apply your filter choices.

To reset filters to the default, click the **Reset** button.

> **Notes:**
> - Filter controls default to **All**.
> - Some filter selections are mutually exclusive. For example, if you select **North America** as the continent, you cannot select individual countries.
> - The **Geography** tab only contains filter controls for **Protocols** and **Segments**.
> - Filter selections you make on one tab are automatically applied to other tabs.
> - When you make filter selections on a tab then navigate to another tab, then navigate back to the original tab, the selections are preserved in the original tab.

Overview Tab

The **Overview** tab serves as a high-level snapshot of your aggregated data and allows you to export data and create Recurring Report Emails and configure Email Alerts.

Summary Area

For various statistics, the Summary Area shows the percent change for the selected reporting date range relative to the prior time period of the same length. For example, if you select **This Month**, the statistics show a comparison between this month's data and the previous month's.

Colors and arrows represent changes:

- An increase displays in green with an arrow pointing up.
- A decrease displays in red with an arrow pointing down.
- If there was no change from the previous period, the text is gray with an empty circle instead of an arrow.
- Percentage up or down is in bold text to the right of the arrow or empty circle.
- Information in gray text beneath the percentage up or down indicates the total date range covered minus any time remaining in the current period. For example, if you select **Last 7 days** as the reporting period and the current date is part of the past seven days, the past 13 days' information is displayed.

> **Note:** Information in the *Summary Area* depends on the selected accounts, time range, time zone, and dimensions.

Statistics in the *Summary Area* are:

Statistic	Description
AVERAGE (Average Throughput)	The average of traffic from the origin, through the network, and out of the network to the requestor. Measured in bps.
(TOTAL) Total Data Transfer	Data from the origin into the CDN and out to the requestor from the published host. Measured in bytes.
(REQUESTS) Number of Requests	Total number of requests from the origin into the CDN and out to the requestor from the published host.

Chart

Data is presented in a chart that shows the total throughput of in bytes plus out bytes across the selected date range by default. The label on the left beneath the summary panel reflects the defaults.

Hover the mouse pointer over the chart to view details in tooltips.

Note:
For any given granularity, if only a single time unit has data, then the data is represented as a single filled circle.

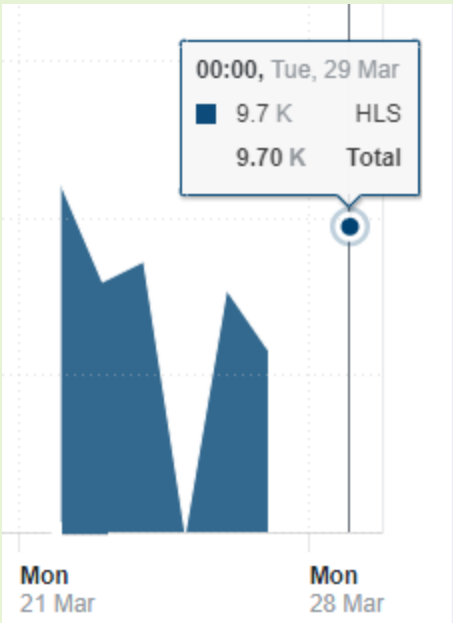

Selecting Chart Granularity

You can refine the chart by selecting a granularity value in the toggle on the left above the chart. Toggles are active depending on the reporting date range (see Selecting a Date Range and Time Zone).

Make a selection in the toggle. The time units on the chart's X-axis are updated to reflect your selection.

Note:
Each value has its own retention policy.

Granularity	Data Retention Policy
5 min	One week
Hour	Five weeks
Day	One year

Note:
The five-minute granularity is not available when:

- The selected time frame is greater than 24 hours.
- When **Continent** or **Country** has been selected in the data grouping drop-down menu.

Selecting a Data Grouping

On the left above the granularity toggle is the report **Grouping** drop-down menu in which you can select up to two options:

- Account
- Protocol
- Segment
- Continent
- Country

Selecting Chart Metrics

On the right beneath the *Summary Area* is the **Metrics**drop-down menu that allows you to select chart metrics. The selected metric is reflected in the y-axis label.

Metric	Description
Throughput	Rate of data flow, measured in bits per second.
Data Transfer	Number of bytes transferred, measured in bytes.
Requests	Number of requests.

You can add an additional dimension to the chosen metrics using the drop-down menu to the right of the metrics drop-down menu.

Selection	Description
In	Ingested into the CDN from the origin.
Out	Bytes from the published host to the end user.
Total	Sum of 'In' and 'Out'.

So for example if you chose 'Throughput' in the metrics drop-down menu and 'In' in the additional drop-down menu, the chart reflect bits per second ingested into the CDN from the origin.

Exporting Chart Data

To export data currently displayed in the chart, click the **Export** drop-down menu on the right above the chart; then choose an option:

- **PowerPoint**: Export to a PowerPoint file that contains a screenshot of the chart.
- **CSV**: Export data to a Comma-Separated Values (CSV) file that reflects the currently selected chart options.

After you choose an option, Control creates and downloads the report.

Toggling Chart Data

A legend below the chart identifies the chart content by color. The legend reflects the choices that you made in the **Grouping**drop-down menu.

For example, if you Chose **Continent** in the drop-down menu, data for continents displays in the chart. Labels for the contents are displayed in the chart legend.

The legend labels are toggles that you can use to display or hide the corresponding chart data. By default, all labels are toggled on. Click a label to hide or show the corresponding data in the chart. Labels that are toggled off have a gray font color.

Creating Recurring Report Emails and Email Alerts

See Working with Recurring Report Emails and Working with Email Alerts in 'Reports General Information'.

Details Tab

The **Details** tab is a deeper drill-down into your data, with additional metrics like “Data transfer IN” or “Data transfer OUT.” Data is presented in a sortable table.

Selecting a Data Grouping

Choose a primary data grouping in the drop-down menu on the left above the table:

- Account
- Protocol
- Segment
- Continent
- Country

Choose a secondary time-related grouping in the drop-down menu on the right above the table:

- Monthly
- Day of week
- Hour of day

> **Note:** When you group by segments, segments are sorted by account and protocol. Master segments are at the top of the list.

Viewing Details in the Table

Each row in the table represents the primary grouping. Use the **[+]** icon to expand rows and reveal details by the time-related grouping.

Adding and Removing Columns

The table loads with three data columns by default, but you can display up to five at a time by adding and deleting columns.

- If fewer than five columns are displayed, add a column by clicking the **+** icon on the right side of the table header row and selecting from the subsequent list.
- Hover over a column and select the **(x)** (Remove column) icon to delete a column.

Each row in the table represents the primary grouping. Use the **[+]** icon to expand rows and reveal details by the time-related grouping.

Exporting Table Data

To export data to a Comma-Separated Values (CSV) file, click the **Export** drop-down menu on the right above the table; then choose a granularity, which determines the content of the exported file.

Granularity	CSV File Contains
Monthly	A total for each metric in each month.
Daily	A total for each metric for each day in each month.
Hourly	A total for each metric for each hour on each day in the month.

After you choose a granularity, Control creates and downloads the report.

CDN Efficiency Tab

The **CDN Efficiency** tab provides additional, more specific information about overall CDN efficiency, expressed as percentages. Delivery configurations and purge events are included to indicate how they influence CDN efficiency. Data is presented in a timeline chart.

Hover the mouse pointer over the chart to view details in tooltips.

Selecting Chart Granularity

You can refine the chart by selecting a granularity value in the toggle on the left above the chart. Each value has its own retention policy.

Granularity	Data Retention Policy
5 min	One week
Hour	Five weeks
Day	One year

Notes:

- For 24 hours or less, the chart offers granularity in five-minute, hourly, and daily increments.
- The five-minute granularity is not available when:
  - The selected time frame is greater than 24 hours.
  - When **Continent** or **Country** has been selected in the data grouping drop-down menu.

Selecting an Efficiency Type

Select a type from the drop-down menu on the right above the table:

- Data Transfer Efficiency
- Requests Efficiency

Toggling Chart Data

On the right side of the chart are toggles that you can use to display or hide the corresponding chart data.

- **Avg**: Average efficiency for the selected type of efficiency.
- **Configs**: Configurations made in Configure > Caching & Delivery or Caching & Delivery (new).
- **Purges**: SmartPurge events.

By default, all labels are toggled on. Click a label to hide or show the corresponding data in the chart. Labels that are toggled off have a gray font color.

> **Note:** Only the configurations and purges relevant to the chosen timeframe are dis

Geography Tab

The **Geography** tab shows traffic for your data categories on an interactive map. The map can be used to understand how your audience traffic varies across the world or your chosen continents or countries.

Hover your mouse pointer over a geographical region to view popups with details about the selected account, time range, timezone, and data breakout, and additional measures.

Selecting a Data Grouping

Choose a data grouping in the drop-down menu on the left above the table:

- Protocol
- Segment
- Account

The selection you make is reflected in the popups that appear when you hover your mouse pointer.

Zooming in and Out

To zoom in on the chart do either of the following:

- Click the **[+]** control on the upper left of the map .
- Click a region (continent, country, or state).

To zoom out, do either of the following:

- Click the **[-]** control on the upper left of the map .
- Click the **Continent** and **Whole World controls** on the right above the chart.

Exporting Data

To export data to a Comma-Separated Values (CSV) file, click the **Export** drop-down menu on the right above the table; then choose a granularity, which determines how data is grouped in the exported file.

Granularity	Grouping in the Exported File
Continents	Data is broken out by continents.

Granularity	Grouping in the Exported File
Countries	Data is broken out by countries within continents.

Hosts & URLs Tab

The **Hosts & URLs** tab shows CDN usage by selected published hostname, URLs, referrer URLs, or file type. You can use the information to analyze content popularity and CDN caching efficiency for individual content URLs. Data is presented in a table.

Values for *% of Total Requests* and *CDN efficiency* include an indicator under the numeric value.

% of Total is a stacked bar chart.

> **Note:**
> URL parameters are stripped from URLs before the metrics are calculated.

Creating Recurring Report Emails

See Working with Recurring Report Emails in 'Reports General Information'.

URLs Displayed

The most requested URLs are displayed, ordered by the '% of Total Requests' column.

By default, the top 50 URLs are displayed based on one-hour samples, but the actual number displayed depends on the timeframe you select in the date picker at the top right of the page. For example, if you select a day as the timeframe, the union of the top 50 URLs of all the one-hour periods in the day would be displayed.

Because the most-requested URLs can differ from one sampling to another, a URL previously displayed might not appear again in the future, or it might appear in a different position in the list.

Data Collection and Availability

Data is not available in realtime; it is processed at the end of each day, and becomes available the next day. As a result, if you select 'Today' or 'Last 24 Hours' in the data range selector, no data is available.

Data is of DAY granularity and is always stored in GMT-7, so the **Time zone**drop-down menu is disabled when you select the tab.

Filtering Table Data

You can limit data in the table to URLs and hostnames that you specify.

- To filter the table, enter all or part of a URL/hostname in the search box on the right above the table and press Enter.
- To display the original list, remove text from the search box and press enter.

Selecting a Data Grouping

Make a selection in the drop-down menu to the right of the search box.

The data displayed in the table depends on the selection; refer to the following table for all data descriptions.

Table Data Item	Description
% of Total Requests	Ratio of the value in the row compared to all rows in the table. To convey the ratio, the column contains: • A horizontal stacked bar chart • The ratio itself
% of Total Traffic	Ratio of the value in the row compared to all rows in the table.
Bytes Efficiency	Total bytes served from cache / (bytes served from cache + bytes served from origin)
CDN Efficiency	How efficient the CDN was in terms of serving end-user requests or bytes from the cache instead of from the origin. Measured in percentage. This column contains a visual indicator divided into four equal quartiles representing 25% each. Each quartile is color-coded: 0 - 25% red 25.xx% – 50% yellow 50.xx% – 75% light green 75.xx% – 100% dark green
Data Transferred	Bytes transferred due to requests for files, broken out by the grouping you selected; for example: bytes transferred by a published host. This is a total of Data coming into the CDN from the origin and data from the published host to the requesting client.
In Bytes	Number of bytes entering the CDN from the origin.
Out Bytes	Total of all byes leaving the CDN through the published hosts.
Out Requests	Number of requests for content leaving the CDN through the published hosts.
Requests	Number of requests for content entering the CDN from the origin.
Total Bytes	'In Bytes' + 'Out Bytes'
Total Requests	'Requests' + 'Out Requests'

Sorting the Table Data

Sort the data in the table by clicking a column header. When you sort on a field, the sort occurs on all pages of data rather than just the current page.

Viewing Row Details

To view a sparkline of average daily traffic for the selected data type (URLs, Published Hosts, and so on), click the row's **[+]** toggle. The toggle changes to **[-]** and the sparkline is displayed.

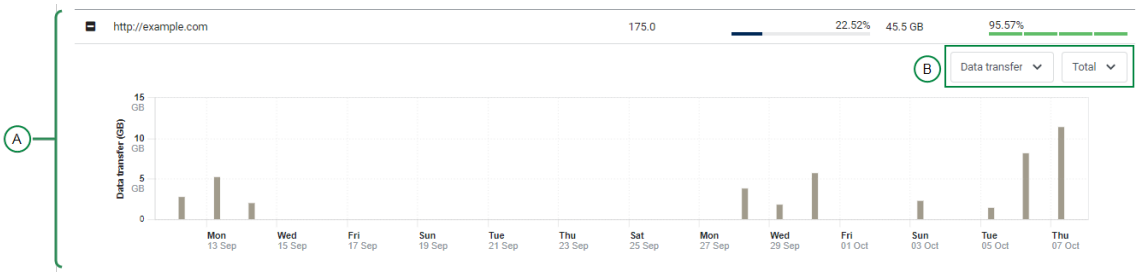

In the preceding figure:

A - Expanded row

B - Metrics drop-down menus

After displaying the sparkline, you can do the following:

- Choose a metric to display using the metrics drop-down menus: 'Requests' , 'Data transfer', 'Total', 'In', 'Out'
- Hover over the sparkline to get additional details in a tooltip.
- Click the **[-]** symbol to hide the sparkline.

Exporting Data

To export chart data to a Comma-Separated Values (CSV), click the **Export** drop-down menu on the right above the table; then select **CSV**.

After you choose **CSV**, Control creates and downloads the report.

> **Note:** Data in the **Hosts & URLs** tab is of DAY granularity and is always stored in GMT-7, so the **Time zone** drop-down menu is disabled when you select the tab.

How Metrics Are Calculated

Metric	Calculation
Throughput In, Out Data Transfer In, Out Requests In, Out	No calculation. Data provided by EdgeQuery.
Throughput Total Data Transfer Total Requests Total	Sum of the corresponding "In" and "Out" values.

Overview Tab - Summary Area

The up or down change for a given metric is calculated as follows.

The metric from the previous period is compared to the metric from the current period.

Comparison Results	Metric Presentation
The selected time period's value equals the previous time period.	Presented as no change.

Comparison Results	Metric Presentation
The previous period's value is smaller than the current period.	Presented as up from the previous period. See also Percentage Calculation.
The previous period's value is greater than the current period.	Presented as down from the previous period. See also Percentage Calculation.

Percentage Calculation

The 'Up' and 'Down' presentations include a percentage.

Percentage = ((newValue - oldValue) / oldValue) * HUNDRED_PERCENT

Content Reports

Content Reports display detail data for content delivered by the CDN. Some of the reports focus on information about requesting users, others on the content delivered.

URL Prefixes Report

The *URL Prefixes Report* shows CDN usage by URL prefix (the host and path portion of the URL), and can be used to analyze content popularity and CDN caching efficiency for individual prefixes.

The report shows the associated *Bytes* (bytes transferred), *Requests* (incoming content requests), or *Seconds* (average connection seconds) for each URL prefix.

Data is aggregated from one or more selected services for the specified date range. The report data can also be filtered by applying a *Segment*, if any are available for the selected *Account Name* and current user.

Data is displayed in both a horizontal bar chart and an interactive data table, and is updated every 1-20 hours.

Up to 15 URL prefixes can be displayed in the bar chart. The displayed URL prefixes correspond to the first 15 URL prefixes shown in the data table.

For each URL prefix, the data table also displays *Ingress* and *Egress* bytes transferred and *CDN Efficiency (%)*.

Status Codes Report

The *Status Codes Report* shows CDN requests by HTTP Status Code and can be used to determine if your origin responds to CDN requests with HTTP error codes.

The report can also show the number of requests by other categories such as services and cache codes.

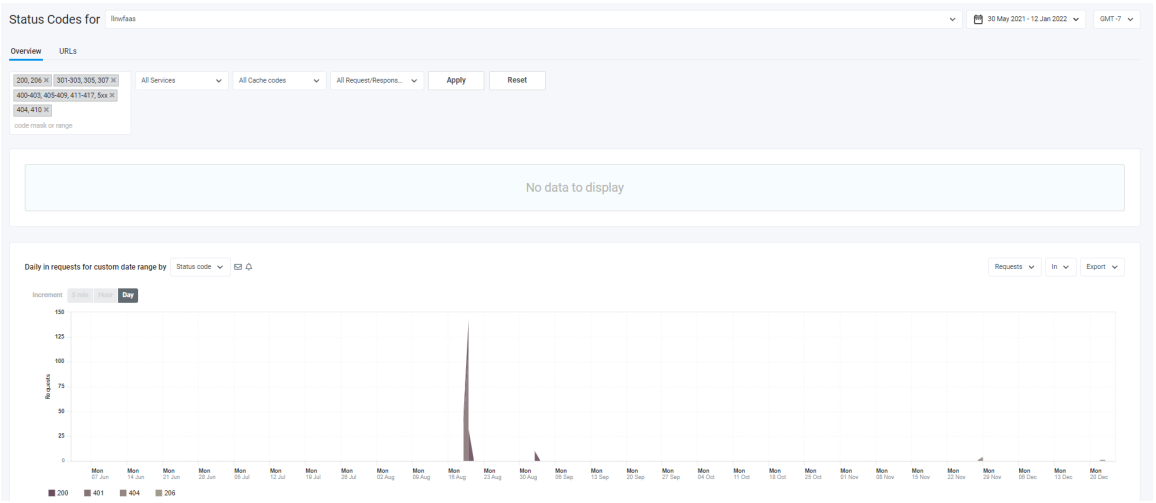

Figure 24. Status Codes Report

The report data is updated every 5 minutes.

The report has these tabs:

Overview

URLs

Report Specifications

Latency	Latency depends on the selected granularity. **Granularity** / **Latency** 5 Minutes / 5 to 10 minutes Hour / 1 hour + delta* Day / 1 day + delta* *delta = approximately 5-10 minutes
Granularity	5 minutes, hour, day
Dimensions	Account, Status Codes, Services, Cache Codes, Request/Response Type
Metrics - Summary Area	Content Served, Redirect, Error, Missing File

Metrics - Elsewhere	Requests (In, Out), Bytes (In, Out),
Delivery Mechanism	EdgeQuery
Associated API Endpoint (s)	• `GET /realtime-reporting-api/traffic/statuscodes` returns status code report data according to the query parameters passed • `POST /realtime-reporting-api/traffic/statuscodes` returns status code report data according to the parameters passed in the request body • `GET /realtime-reporting-api/traffic/statuscodes/cachecodes` returns a list of cache codes • `GET /realtime-reporting-api/traffic/statuscodes/requestresponsetype` returns a list of request response types • `GET /realtime-reporting-api/traffic/statuscodes/retentions` returns a list of retentions • `GET /realtime-reporting-api/traffic/statuscodes/services` returns a list of services

Selecting an Account

Select one or more accounts in the **Status Codes for**drop-down menu at the top of the page.

Notes:
- If you don't select an account, the accounts for the default company are automatically selected, and a warning is displayed.
- If you modify the selections in the **STATUS CODES FOR** drop-down menu, the filter selections you made are reset to defaults and the chart updates automatically to match the defaults.

Selecting a Date Range and Time Zone

Select a date range and timezone in the drop-down menus at the upper right part of the page.

The selected date range influences the values on the chart x-axis.

Date Range	Increment
Three days or less	5-minute data increments
Greater than three days and less than or equal to one month	1-hour data increments
Larger than one month	1-day data increments are displayed.

Note: Data in the **URLs** tab is of DAY granularity and is always stored in GMT-7, so the **Time zone**drop-down menu is disabled when you select the tab.

Overview Tab

Summary Area

Summary Area

For various statistics, the Summary Area shows the percent change for the selected reporting date range relative to the prior time period of the same length. For example, if you select **This Month**, the statistics show a comparison between this month's data and the previous month's.

Colors and arrows represent changes:

- An increase displays in green with an arrow pointing up.
- A decrease displays in red with an arrow pointing down.
- If there was no change from the previous period, the text is gray with an empty circle instead of an arrow.
- Percentage up or down is in bold text to the right of the arrow or empty circle.
- Information in gray text beneath the percentage up or down indicates the total date range covered minus any time remaining in the current period. For example, if you select **Last 7 days** as the reporting period and the current date is part of the past seven days, the past 13 days' information is displayed.

> **Note:** Information in the *Summary Area* depends on the selected accounts, time range, and time zone.

Statistics in the *Summary Area* are:

Statistic	Description
CONTENT SERVED	Content was served to the requestor.
REDIRECT	Requests were redirected.
ERROR	Requests resulted in an error condition.
MISSING FILE	Requested file not found.

Chart

Filtering Chart Data

Using filter controls at the top of the report, you can filter on **Status Codes**, **Services**, **Cache Codes**, and **Cache Conditions**.

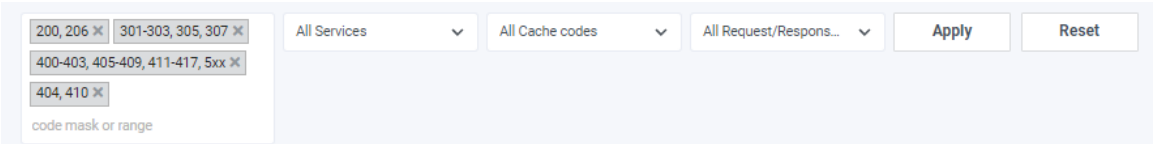

Make the desired selections, then click the **Apply** button to apply your filter choices.

To reset filters to the default, click the **Reset** button.

> **Notes:**
> - Filter controls default to **All**.
> - If you modify the selected accounts in the **Status Codes for** drop-down menu (see Selecting an Account), the filter selections you made are reset to defaults, and the chart is automatically updated to match the defaults.
> - When you change the status codes filter, navigate to the **Overview Tab**, then navigate back to the **URLs Tab**, the filter selections are preserved in the **URLs Tab**.

'Filter by' Selections

Selection	Additional Control	Instructions
Status codes	A text field to configure desired status codes.	Do any of the following: • Click in the field and select pre-configured status code ranges. • Enter your own status code ranges. • Enter a single status code. **Note:** Duplicates status codes and ranges are not allowed.
Services	A drop-down menu of services.	Select one or more values from the drop-down menu. **Note:** The selections available depend on the services for which your company has signed up.
By Cache codes	A drop-down menu of cache codes.	Select one or more values from the drop-down menu: • **Miss** - cache miss; response is delivered from customer origin • **Redirect** - request redirected • **Hit** - cache hit; response is delivered from cache • **Other -** not one of the above
By Request/Response type	A drop-down menu of request and response types.	Select one or more values from the drop-down menu: • **Standard** - not any of the below • **If modified since** - client-issued 'If-Modified-Since' requests only • **Negatively cached** - caching of 4xx and 5xx responses • **Refresh** - CDN-issued refresh checks only

Selecting Chart Granularity

You can refine the chart by selecting a granularity value in the toggle on the left above the chart. Toggles are active depending on the reporting date range (see Selecting a Date Range and Time Zone).

Make a selection in the toggle. The time units on the chart's X-axis are updated to reflect your selection.

Note:
Each value has its own retention policy.

Granularity	Data Retention Policy
5 min	One week
Hour	Five weeks
Day	One year

Selecting a Data Grouping

On the left above the granularity toggle is the report grouping drop-down menu in which you select a chart grouping.

- Status codes
- Accounts
- Service
- Cache code
- Request/Response type

Make a selection. The chart content and labels beneath the chart are updated to reflect the selected option.

Note: When you select **Account**, the accounts displayed reflect the selected accounts in the **Status Codes for**drop-down menu (see Selecting an Account). All other selections reflect the chart filters applied to the chart (see Filtering Chart Data).

Selecting Chart Metrics

Toggling Between Requests and Bytes

By default, the chart displays the number of requests as the value for the y-axis. Users with appropriate permissions can also choose to view bytes instead of the number of requests using a drop-down menu on the right side of the screen under the summary bar:

- All customers see 'Requests' in the drop-down menu.
- Customers with additional permissions also see a second entry, 'Bytes'.

Customers who see both values can toggle between the two:

1. Select a value from the **Requests/Bytes**drop-down menu under the summary bar on the right side:
   - Requests - display requests (default)
   - Bytes - display bytes
2. The chart refreshes to reflect the selections you made.

Note: Contact your account manager to learn more about the 'Bytes' option.

Toggling Between Outgoing or Incoming Requests

1. Select a value from the **Out/In**drop-down menu under the summary bar on the right side:
   - Out - Outgoing requests (default)
   - In - Incoming content requests
2. The chart refreshes to reflect the selections you made.

Toggling Chart Data

A legend below the chart identifies the chart content by color. The labels reflect the data grouping (see Selecting a Data Grouping) in the chart.

The legend labels are toggles that you can use to display or hide the corresponding chart data. By default, all labels are toggled on. Labels that are toggled off have a gray font color.

Click a label to hide or show the corresponding data in the chart. The chart content is updated to reflect your selection.

Exporting Chart Data

To export data currently displayed in the chart, click the **Export** drop-down menu on the right above the chart; then choose an option:

- **PowerPoint**: Export to a PowerPoint file that contains a screenshot of the chart.
- **CSV**: Export data to a Comma-Separated Values (CSV) file that reflects the currently selected chart options.

After you choose an option, Control creates and downloads the report.

Creating Recurring Report Emails and Email Alerts

See Working with Recurring Report Emails and Working with Email Alerts in 'Reports General Information'.

> **Note:** You can also create email alerts in the URLs tab. For more information, see Configuring Email Alerts Per URL.

URLs Tab

Overview

The URLs tab shows:

- Missing Files information: requested files that are missing on your origin server and can be used to find and fix bad links and omitted files. A file is determined to be missing if your origin server returned an HTTP 404 status code during initial cache fill or a subsequent freshness check by the CDN.
- File Errors information: requested files for which your origin server returned an HTTP error code and can be used to find and fix problems with specific files. A file is determined to have an error if your origin server returned an HTTP error code during initial cache fill or a subsequent freshness check by the CDN.
  Any 5xx response is considered an error and appears in the report.

Notes:

- Data in the **URLs** tab is of DAY granularity and is always stored in GMT-7, so the **Time zone**drop-down menu is disabled when you select the tab.
- Data is not available in realtime; it is processed at the end of each day, and becomes available the next day. As a result, if you select 'Today' or 'Last 24 Hours' in the data range selector, no data is available.
- The number of reported URLs is limited to 50.

The tab shows data in a table with three columns. You can sort the data by any column.

- **URLs**: File URLs
- **Requests**: Number of requests for the file that resulted in an error
- **% of Total**: Percent of requests for the file out of all requests for all files. Each row has a horizontal bar chart representing the percentage.

Data is aggregated by service (HTTP, HTTPS) for the specified date range.

Selecting an Error Type

Select 'Missing Files' or 'File Errors' in the dropdown under the tab header on the page's right side.

Viewing Error Details

1. Click the **+** icon on a row to expand the row and display a chart.
   The chart shows the number of requests per time unit in the selected date range. The time unit depends on the date range you selected.
2. Move your mouse over the chart to see details for any given time unit.
3. Click the **-** icon to hide the chart.

Filtering the URL List

You can filter data either by URL or status codes.

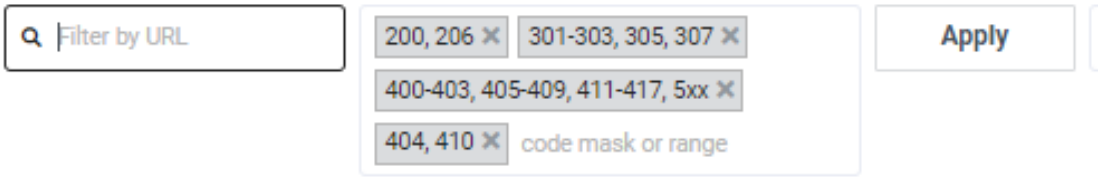

- To filter by URL, type a phrase in the search field, then press the **Enter** key on your keyboard.
- To filter by status codes, add, modify, or delete single status codes, or ranges of codes, then click the **Apply** button.

After the filter is instated, the contents of the table change to reflect your filter.

When you change the status codes filter, navigate to the **URLs Tab**, then navigate back to the **Overview Tab**, the filter selections are preserved in the **Overview Tab**.

Sorting Data

Click a column heading to sort data by that column. A bar is displayed above or below the column heading to indicate if the data is sorted in ascending or descending order, respectively.

Configuring Email Alerts per URL

You can configure email alerts on a URL basis, per your published host domain. See Working with Email Alerts in 'Reports General Information'.

Exporting URLs Data

To export data to a Comma-Separated Values (CSV) file, click the **Export** drop-down menu to the right above the chart; then select **CSV**.

After you select **CSV**, Control creates and downloads the report.

How Metrics Are Calculated

Metric	Calculation
Requests In, Out Bytes In, Out	No calculation. Data provided by EdgeQuery.

Overview Tab - Summary Area

The up or down change for a given metric is calculated as follows.

The metric from the previous period is compared to the metric from the current period.

Comparison Results	Metric Presentation
The selected time period's value equals the previous time period.	Presented as no change.
The previous period's value is smaller than the current period.	Presented as up from the previous period. See also Percentage Calculation.
The previous period's value is greater than the current period.	Presented as down from the previous period. See also Percentage Calculation.

Percentage Calculation

The 'Up' and 'Down' presentations include a percentage.

Percentage = ((newValue - oldValue) / oldValue) * HUNDRED_PERCENT

Realtime Live Event Overview Report

The *Realtime Live Event Overview Report* shows a variety of metrics for live *HLS*, *HDS* and *MSS* streams, and can be used to analyze attendance and audience behavior during the events.

Report data is shown for the selected services (*HLS*, *HDS* and/or *MSS*).

The report includes a *Live Events Summary* table for "top events", a *Live Events Summary* line chart that displays audience size for the selected events for the preceding 6 hours, and a *Global Bitrate Distribution* vertical bar chart that displays audience size by event bitrate.

An *Events* table provides detailed information for each event, including *Concurrent Views*, *Length of View*, *% HLS*, *% HDS, % MSS* and a link to a *Live Event Details* page.

The report also includes an interactive map (similar to the Geography tab of the *Traffic Report*) that displays aggregate audience size by geographic area, and a *Live Event Geo Summary* data table that shows audience distribution by major geographic area (NA, APAC, SA and EU) and by bitrate.

Data is updated every 120 seconds.

Storage Reports

Storage reports provide usage metrics for Limelight CDN storage services.

Origin Storage Report

The Origin Storage Report lets Origin Storage (previously *Cloud Storage*) users view consolidated storage data for any combination of accounts and policies in a single, comprehensive report. Users with permissions for the report can view a summary of stored content across time, peak storage, total stored on disk and total unique objects.

Data is displayed in an interactive area chart.

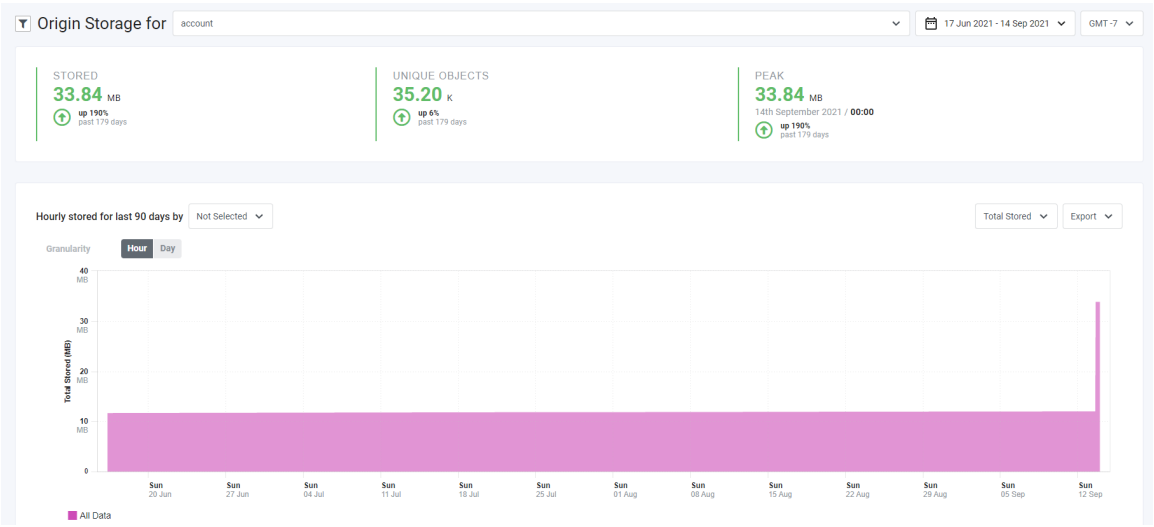

Latency	Latency depends on the selected granularity. **Granularity** / **Latency** Hour / 1 hour + delta Day / 1 day + delta delta = approximately 5-10 minutes
Granularity	Hour, day
Dimensions	Policy, Account
Metrics - Summary Area	Total Storage in bytes, Unique Objects, Peak Storage in bytes
Delivery Mechanism	EdgeQuery
Associated API Endpoint	None

(s)	

Selecting Accounts, Date Range, and Time Zones

You can make selections in the controls above the tab header:

- **ORIGIN STORAGE FOR**. Select one or more accounts to which you have access for cross-account analysis. Click the **Select All** button to select all accounts.

  > **Note:** You must select at least one account; otherwise, the default company is automatically selected, and a warning is displayed.

- **Date range**. Pick from pre-set time frames or choose custom date ranges in the drop-down menu. Click the **Apply** button on custom ranges.
- **Time zone**. The top five most commonly used timezones in Control are at the top of the drop-down menu. Scroll down for additional time zones.

  > **Note:** Data in the **Hosts & URLs** tab is of DAY granularity and is always stored in GMT-7, so the **Time zone** drop-down menu is disabled when you select the tab.

Filtering

A filter icon on the top left side of the page allows you to filter by policy. The icon is a toggle that displays and hides the filter controls.

1. Click the icon to display a drop-down menu.
2. Select **Policy** in the drop-down menu to display a drop-down menu with policy names related to the selected accounts.
3. Select the desired policies, then click the **Apply** button to apply your filter choices.

Your filter choices are applied to the *Summary Area* and area chart.

Summary Area

For various statistics, the *Summary Area* shows the percent change for the selected reporting date range relative to the previous time range of the same duration. Colors and arrows represent changes:

- An increase displays in green with an arrow pointing up.
- A decrease displays in red with an arrow pointing down.
- If there was no change from the previous period, the text is gray with an empty circle instead of an arrow.
- Percentage up or down is in bold text to the right of the arrow or circle.
- Information in gray text beneath percentage up or down indicates the total date range covered minus any time remaining in the current period. For example, if you selected **Last 7 days** as the reporting period and the current date is part of the past seven days, the information displayed is for the past 13 days.

> **Note:** Information in the *Summary Area* is determined only by the selected account and reporting date range. It is not affected by the selections you make in any other controls on the page.

Statistics in the *Summary Area* show these values for selected accounts, time range, time zone and policies:

- *STORED*: total size in bytes of all objects stored.
- *UNIQUE OBJECTS*: Number of objects stored in Origin Storage at any given moment of time, not including copies. According to some storage policies , Limelight maintains several copies. The Unique Objects value does not include copies in the count.
- *PEAK*: Date and time with the highest number of bytes stored.

Selecting a Data Grouping

Beneath the *Summary Area* title is a drop-down menu that allows you to determine how data is broken out in the chart, by Policy, Account or all data.

Make a selection.

Selection	Data Displayed in Chart	Information Displayed in Legend
Not Selected	All data for the selected metric.	Single entry: **All Data**.
Policy	Data for the filtered policies.	List of filtered policies.
Account	Data for all selected accounts.	List of selected accounts.

The chart reflects your selection.

Choosing Metrics

On the right beneath the *Summary Area* is a drop-down menu that allows you to select chart metrics.

Make a selection.

Selection	Description
Total Stored	Total size in bytes of all objects stored.
Unique Objects	Number of objects stored in Origin Storage at any given moment of time, not including copies. According to some storage policies , Limelight maintains several copies. The Unique Objects value does not include copies in the count.

The chart reflects your selection.

Choosing Granularity

The chart can be further refined by selecting one of the **Increment** values:

- Hour
- Day

The selection you make determines increments along the Y-axis.

Exporting Chart Data

To export chart data to a PowerPoint file that contains a screenshot of the chart, click the **Export**drop-down menu on the right above the chart; then select **PowerPoint**.

After you select **PowerPoint**; Control creates and downloads the report.

Toggling Chart Data

A legend beneath the chart identifies the chart content by color. The legend reflects the choices that you made in the policy filter and the grouping drop-down menu.

The legend labels are toggles that you can use to display or hide the corresponding chart data. By default, all labels are toggled on. Click a label to hide or show the corresponding data in the chart. Labels that are toggled off have a gray font color.

How Metrics Are Calculated

Metric	Calculation
Total of stored objects' sizes	No calculation. Data provided by EdgeQuery.
Unique objects	Total number of objects minus copies.
Peak storage	No calculation. Data provided by EdgeQuery.

Summary Area - Percentage Up and Down

The up or down change for a given metric is calculated by comparing the metric from the previous period to the metric from the current period.

Comparison Results	Metric Presentation
Selected time period value equals previous time period.	Presented as no change.
Previous period value is smaller than current period.	Presented as up from previous period. See also Percentage Calculation.
Previous period value is greater than current period.	Presented as down from previous period. See also Percentage Calculation.

Percentage Calculation

The 'Up' and 'Down' presentations include a percentage.

Percentage = ((newValue - oldValue) / oldValue) * HUNDRED_PERCENT

Transit Report

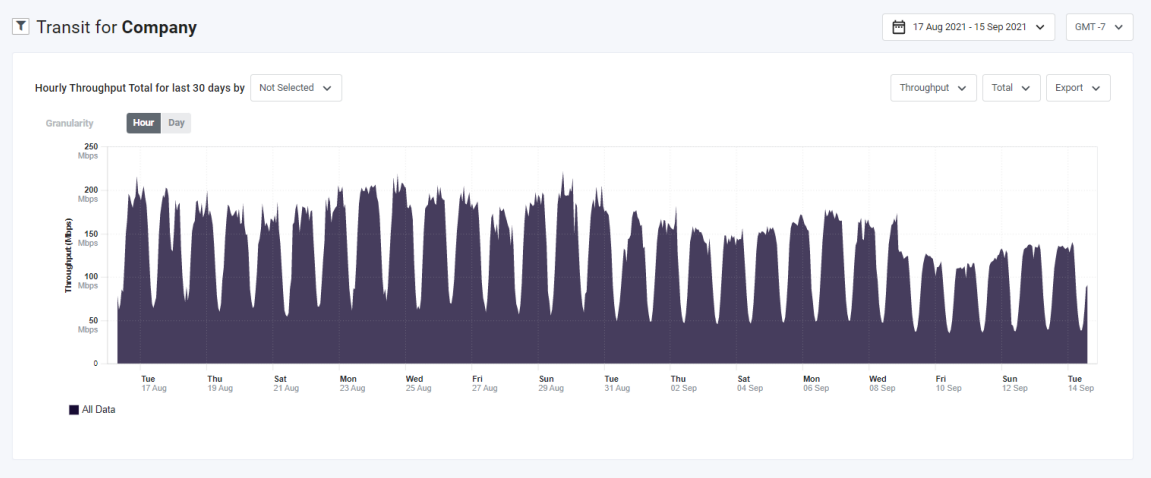

The Transit Report shows CDN transit usage (access to the Limelight CDN backbone) for the currently selected company. You can use the report to monitor traffic trends for the IP Connect service. The report shows throughput and data transfer.

> **Note:** Backbones and circuits are the private infrastructure that connect Limelight's PoPs.

Data is displayed in an interactive chart.

Report Specifications

Latency	Depends on the selected granularity. • Hour granularity: latency = 1 hour + Delta • Day granularity: latency = 1 day + Delta Delta = approximately 5-10 minutes
Granularity	Hour, Day
Dimensions	Circuit
Metrics	• Throughput in, throughput out, and throughput total (in + out). • Data transfer in, data transfer out, and data transfer total (in + out). See How Metrics Are Calculated for details.
Delivery Mechanism	Realtime Reporting API
Associated API Endpoint	`http://{host}/realtime-reporting-api/transit` Returns transit data for the requested time period.

Choosing Date Range and Time Zone

Use the controls at the top right of the page:

- **Date Range Control**. Pick from pre-set time frames or choose custom date ranges in the drop-down menu. Click the **Apply** button to set a custom range.
- **Time Zone Control**. Select a timezone. The five most commonly used timezones in Control are at the top of the drop-down menu. Scroll down for additional time zones.

Choosing Circuits to Display

Use the **Circuit Name Filter** and the **Grouping Drop-Down Menu** to determine circuit data to display.

Circuit Name Filter

The filter icon (a funnel) on the left side of the *TRANSIT FOR* title allows you to filter by circuit name. The icon is a toggle that displays and hides the circuit filter control.

1. Click the filter icon.
2. Make a selection in the subsequent **Filter by**drop-down menu.
   By default, the **Circuit Name** entry is unchecked. As such, data for all circuits is displayed in the chart.
3. If desired, select **Circuit Name** to display a drop-down menu with circuit names you can select to display in the chart.
4. Click the **Apply** button.

Grouping Drop-Down Menu

Beneath the *TRANSIT FOR* title is a drop-down menu that allows you to display all circuit data or data for specific circuit names.

Make a selection:

- Choose **Not Selected** to display all data. The legend displays a single entry, **All Data**.
- Select **Circuit** to display data broken out by circuit name in the legend.

Choosing Chart Metrics

Beneath the **Date Range Control** and **Timezone Control** are drop-down menus that allow you to select report metrics.

Throughput/Data Transfer

Make a selection:

- **Throughput** - Data flow rate through the network expressed in Kbps.
- **Data Transfer** - Amount of data transferred expressed in GB.

In/Out/Total

This is a second-level qualifier for the Throughput/Data Transfer. Make a selection:

- **In** - From origin to the CDN.
- **Out** - From the CDN to the end user.
- **Total** - Total of In + Out

For example, if you choose **Data Transfer** and **Total**, the chart displays total (transfer in + transfer out).

Setting Chart Date Granularity

The chart X-axis displays date/time units, which are determined by the **Granularity Controls** and the **Date Range Control**.

Granularity Controls

Granularity Controls are located above the chart on the left. The selection you make determines the data available in the chart.

- **Hour** - Display hourly data.
- **Day** - Display only daily data.

Date Range Control

You can select ranges as small as a day or ranges spanning multiple months.

Scenario 1: You select a single day or last twenty-four hours.

Hour Granularity	Day Granularity
Hours for the day are displayed on the X-axis. Move your cursor across the chart to see metrics for each hour.	One data point representing midnight on the selected day is displayed on the X-axis. Move your cursor over the chart to see metrics for the date.

Scenario 2: You select a date range that includes more than one day.

Hour Granularity	Day Granularity
Dates are displayed on the X-axis. Move your cursor across the chart to see metrics for each hour in the dates.	Data points representing dates at midnight are displayed on the X-axis. Move your cursor across the chart to see metrics for each date.

Note:
For 'Day' granularity, the dates displayed depend on the size of the date range. The larger the range, the fewer the dates displayed, but data for all dates can be viewed. For example, for smaller date ranges, a data point for each date is displayed. For larger ranges, data points for every seven days might be displayed, but by moving the cursor across the chart, you can see popups with data for each day.

Toggling Chart Data

A legend beneath the chart identifies the chart content by color. The legend reflects the choices that you made in the **Circuit Name Filter** and the **Grouping Drop-Down Menu**.

For example, if you filtered by two specific circuit names and you selected Circuit in the drop-down menu, data for the circuit names displays in the chart. Labels for the two names display in the chart legend.

The legend labels are toggles that you can use to display or hide the corresponding chart data. By default, all labels are toggled on. Click a label to hide or show the corresponding data in the chart. Labels that are toggled off have a gray font color.

Note: If you chose **Not Selected** in the **Grouping Drop-Down Menu**, the label **All Data** is displayed. Toggling the label hides or reveals all chart data.

Exporting Chart Data

To export data currently displayed in the chart, click the **Export** drop-down menu on the right above the chart; then choose an option:

- **PowerPoint**: Export to a PowerPoint file that contains a screenshot of the chart.
- **CSV**: Export data to a Comma-Separated Values (CSV) file that reflects the currently selected chart options.

After you choose an option, Control creates and downloads the report.

How Metrics Are Calculated

Metric	Calculation
Data transfer in	No calculation. Data provided by EdgeQuery
Data transfer out	No calculation. Data provided by EdgeQuery
Data transfer total	Data transfer in + Data transfer out
Throughput in	"Data transfer in" divided by "duration," where duration depends on the chosen granularity. • Day granularity: For each day, the system divides Data transfer in by the number of seconds in the day. • Hour granularity: For each hour, the system divides Data transfer in by the number of seconds in an hour. Hover the mouse pointer over a day or hour to display the results of the calculation.
Throughput out	Calculation is identical to 'Throughput in', but applied to Data transfer out.
Throughput total	Throughput in + Throughput out