

Edgecast

Analytics Suite User Guide

edgecast

Disclaimer

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

Trademark Information

EDGECAST is a registered trademark of Edgecast Inc.

About This Guide

Analytics Suite User Guide

Version 3.30

11/10/2022

© 2022 Edgecast Inc. All rights reserved.

Table of Contents

Analytics Suite.....	1
Overview	1
Accessing Reports	2
Viewing/Generating Reports	2
Filtering by POP.....	3
Filtering by Time Period	3
Reporting CDN Activity	6
HTTP Streaming.....	6
Unit Conversion	6
Exporting Report Data	7
Core Reporting Module	8
Overview	8
Logged Data	8
Platform-Specific Reports	8
Traffic Summary.....	9
Bandwidth.....	11
Data Transferred.....	13
Hits (Status Codes).....	16
Cache Statuses	17
Cache Hit Ratio.....	18
Cnames (Deprecated)	20
CDN Storage Usage	21
Edge Image Optimizer Usage	22
DNS.....	23
Custom Reports.....	24
Overview	24

Logging	24
Edge CNAMEs.....	25
Advanced Content Analytics	27
Overview	27
Reports & Log Collection.....	27
Asset/Directory Location	27
Advanced HTTP and ADN Reports	28
Geography Reports (Map-Based).....	28
Geography Reports (Bar Charts)	30
Daily Summary Report	32
By Hour Report.....	33
By File Report	34
By File Detail Report.....	35
By File Type Report	36
By Directory Report.....	37
By Browser Report	38
By Referrer Report	39
By Download Report	40
By 404 Errors Report.....	42
Real-Time Statistics	43
Overview	43
Overview Report	43
Detailed Real-Time Statistics	47
Other Stats	53
Edge Performance Analytics	54
Overview	54
Reports & Log Collection.....	54
Asset/Directory Location	54
Dashboard.....	55
Chart.....	56
Timeline.....	57
Key Metrics & Statistics.....	58

Scheduled E-mails and Alerts.....	76
Scheduled Report Delivery.....	76
Alerts.....	77
Reports.....	80
Daily Summary Report.....	80
Hourly Summary Report.....	81
Protocols Report.....	82
HTTP Methods Report.....	83
URLs Report.....	84
Cnames Report.....	85
Origins Report.....	86
Geo POPs Report.....	87
Clients Report.....	88
Cache Statuses Report.....	89
NONE Details Report.....	90
CONFIG_NOCACHE Details Report.....	91
UNCACHEABLE Details Report.....	92
TCP_HIT Details Report.....	93
TCP_MISS Details Report.....	94
TCP_EXPIRED_HIT Details Report.....	95
TCP_EXPIRED_MISS Details Report.....	96
TCP_CLIENT_REFRESH_MISS_Details Report.....	97
Client Request Types Report.....	97
User Agents Report.....	99
Referrers Report.....	100
Compression Types Report.....	100
File Types Report.....	102
Unique Files Report.....	103
Token Auth Summary Report.....	104
Token Auth Deny Details Report.....	106
HTTP Response Codes.....	107
400 Errors.....	107

401 Errors.....	108
404 Errors.....	108
403 Errors.....	109
429 Errors.....	110
4xx Errors	110
504 Errors.....	111
503 Errors.....	112
502 Errors.....	113
5xx Errors	114
Real-Time Alerts.....	115
Overview	115
Alert Criteria.....	115
Notifications.....	116
E-mail Notifications.....	116
HTTP (Post) Notifications	117
Notification Keywords.....	118
Real-Time Alert Administration	120
Creating a Real-Time Alert	120
Modifying a Real-Time Alert	122
Deleting a Real-Time Alert	122
Finding a Real-Time Alert	122
User Interface	123
Alert Configurations.....	123
Appendix A.....	128
Custom Report Fields.....	128
Advanced Content Analytics Fields.....	130
Edge Performance Analytics Fields	133
Common Edge Performance Analytics Fields	133
Edge Performance Analytics Fields	135
Appendix B.....	139
Monitoring Criteria (Metrics).....	139
Appendix C.....	144

Cache Statuses	144
Glossary.....	146

Analytics Suite

Overview

The Analytics Suite consists of the following modules:

Module	Purchase Separately?	Data Availability	Update Frequency
Core Reporting	No	18 months	Continuous
Custom Reports	No	18 months	Continuous
Advanced Content Analytics	Yes	90 days	Daily (Previous Day)
Real-Time Statistics	Yes	Real-time	Real-time
Edge Performance Analytics	Yes	90 days	Daily (Previous Day)
Real-Time Log Delivery	Yes	Near real-time	Near real-time
Report Builder	Yes	90 days	2 minutes

Use these reports to check traffic activity, bandwidth usage, storage usage, and cache activity statistics. These reports provide insight into how our CDN is delivering data to your clients. The information gleaned from these reports will help you analyze data usage patterns.

Note: The statistical information used to generate these reports is mined from data logged by our servers. A backup version of CDN activity data can be stored on either our CDN origin server or on a server of your own choosing. For more information, please refer to the **Log Files** chapter.

Accessing Reports

All available modules are listed under the **Analytics** menu.

Reminder: If a link is disabled, then that reporting module has not been purchased. Please contact your CDN account manager for additional details.

Once you have selected the desired reporting module, a side navigation bar will display the available reports for the selected module. When navigating through the available reports, keep the following in mind:

- **Solid Arrow (▶):** A solid arrow indicates that there are additional types of reports available. Typically, clicking a solid arrow will result in additional sub-reports appearing directly below it. The expanding and collapsing nature of this report section is indicated by a solid down arrow (▼). If selecting a solid arrow doesn't expand/collapse a section, then an additional report can be viewed from within the report that is displayed for that item.
- **Hollow Circle (○):** A hollow arrow indicates that a report will be displayed when you select that item.

Viewing/Generating Reports

A report can be viewed by simply navigating to the desired reporting module (i.e., Core Reporting, Advanced Content Analytics, Real-Time Statistics, or Edge Performance Analytics) and then clicking on the desired report.

Note: The majority of the reports in the Core Reporting module also require that you select the platform on which a report will be generated.

Once the desired report page has loaded, you can choose to filter the data that will be used to generate the report. The available filtering options depend on the type of report being viewed. Most reports allow you to filter by POP and/or time period. Both of these options are explained below.

Note: There are no filtering options for Real-Time Statistics reports. Real-Time Statistics always display current information on CDN activity for your account.

Filtering by POP

A common filtering option for the Core Reporting module is **Edge Nodes** (i.e., POPs). This option allows you to filter for data from a particular POP. If you would like to see data across the entire network, then you will need to ensure that the "All Edge Nodes" item is selected from the **Edge Nodes** option.

Note: The data displayed in the currently selected report will be automatically updated when you select a POP from the **Edge Nodes** option.

Filtering by Time Period

Most reports allow you to select the time period for which a report will be generated. You can either select a date range using a predefined time period or you can specify a custom date range. The available filtering options for selecting a time period are described below.

Time Period	Description
<i>Month Year</i>	Indicates that the report will display data for the selected month/year combination (e.g., July 2015). This period starts at midnight on the first day of the month and ends at 23:59 on the last day of the month.
Past 24 Hours	Indicates that the report will display data for the last 24 hours. After you select this option, the starting date and ending date options will be updated to indicate a 24 hour period that ends at the top of the hour of the current hour. For example, if it is currently 10:30 a.m., then the 24 hour period would start at 10 a.m. on the previous day and end at 10 a.m. on the current day.
Today	Indicates that the report will display data for the current day. After you select this option, the starting date/time will be updated to read 00:00 (i.e., midnight) of the current day. The end date/time will be the current date and time.
This Week	Indicates that the report will display data for the current week. After you select this option, the starting date/time will be updated to read 00:00 (i.e., midnight) for a date 7 days prior to today. The end date/time will be the current date and time.
This Month	Indicates that the report will display data for the current month. After you select this option, the starting date/time will be updated to read 00:00 (i.e., midnight) for the first day of the current month. The end date/time will be the current date and time.
Yesterday	Indicates that the report will display data for yesterday. After you select this option, the starting date/time will be updated to read 00:00 (i.e., midnight) of the previous day. The end date/time will be 23:59 of the previous day.

Time Period	Description
Last Week	Indicates that the report will display data for the last week. After you select this option, the starting date/time will be updated to read 00:00 (i.e., midnight) for a date that is 14 days prior to the current date. The end date/time will be 23:59 for a date 7 days prior to the current date.
Last Month	Indicates that the report will display data for the previous month. After you select this option, the starting date/time will be updated to read 00:00 (i.e., midnight) for the 1 st of the previous month (e.g., 2015-07-01 00:00). The end date/time will be 23:59 for the last day in that month (e.g., 2015-07-31 23:59).
Custom	<p>Selecting this item allows you to choose the starting and ending date/time for your report. You will be able to choose the starting and ending date/time from the From and the To option, respectively.</p> <hr/> <p>Tip: If you choose to specify a custom date range, then make sure that you limit it to approximately 1 month. Specifying a longer time frame may cause a long delay or even prevent the report from being generated.</p> <hr/>

Note: The type of options that will be available for time period varies by report type.

Note: The data displayed in the currently selected report will be automatically updated when you select a predefined time period from the **Date Range** option. However, if you choose the **Custom** item, then you will need to manually click **Go** before the report will be updated.

Relationship between Start/End Time and Data Reported

As noted previously, the time period used to generate a report plays an important role in determining the bulk of the CDN activity data that will be reported. However, there is another factor that determines whether additional CDN activity will be included in the reported value. This factor is that start and end date/times are inclusive. In order to understand what this means, you will need to know that data is reported in chunks of time (e.g., 5 minutes, 1 hour, 1 day, etc.). A report will include all of the chunks that fall within the specified time period and the chunks that correspond to the specified start and end date/time. This will occur regardless of whether the specified start and end date/time falls at the start, middle, or end of a chunk of time.

A few key facts about how time chunks affect or interact with report data:

- The amount of time covered by a chunk varies for each type of report. This time interval can either be 5 minutes, 1 hour, 1 day, or 1 month.
- A report's time chunk should not be confused with the date/time range used to generate the report. Please refer to the documentation provided for the desired report to find out the time chunk that it uses to report data.

- If a start/end time cannot be specified for a report, then the report will include data for the specified start/end date. For example, specifying a date range of "2015-08-14 to 2015-08-15" will include data for both 8/14/2015 and 8/15/2015.
- The start and end time specified for Edge Performance Analytics reports is irrelevant. It will always include all data that took place on the specified start and end date (as demonstrated in the above note).
- A start or end date/time cannot be specified for monthly reports (e.g., Traffic Summary). A date range is not displayed for this type of report, since report data will always be limited to the specified month (e.g., 08-01-2015 00:00:00 – 08-31-2015 23:59:59 GMT). As a result, this case is not covered in this section.

It is important to know the following information when generating or viewing a report:

- What is the report's start and end date/time?
- What type of time chunk is used to report data?

The above information can be used to identify the exact time period that will be covered by a report. Simply make sure to account for the chunk of time used by the report to figure out the exact time period for which CDN activity will be reported.

The following table illustrates how the exact time period that will be included in a report can be calculated for each type of time chunk. This example assumes that the following date/time range was used to generate the report:

- **Start Date/Time:** 2015-08-01 07:02:00
- **End Date/Time:** 2015-08-02 00:00:00

Time Chunk	Actual Start Date/Time	Actual End Date/Time
5 Minutes	2015-08-01 07:00:00	2015-08-02 00:04:59
1 Hour	2015-08-01 07:00:00	2015-08-02 00:59:59
1 Day	2015-08-01 00:00:00	2015-08-02 23:59:59

Reporting CDN Activity

Several reports (e.g., Traffic Summary, Bandwidth, and Data Transferred) provide information on the CDN response to the client. The data used to calculate these reports is generated from the CDN response that is sent from a POP to a client.

Note: If the requested content is protected by the Origin Shield feature, then it will also include the traffic that is sent from the Origin Shield server to the POP serving the request.

Note: If the requested report is for the ADN platform, then it will also include traffic from the ADN gateway server to the POP serving the request.

HTTP Streaming

We offer the following HTTP streaming solutions: HTTP Progressive Download, Smooth Streaming, Dynamic Cloud Packaging, and Media Ingest. Traffic statistics for all of these HTTP streaming solutions are included in reports generated for the HTTP Large platform.

Unit Conversion

Although CDN activity data is stored in bytes, the value displayed in a report is a much larger unit (e.g., Gigabytes or Terabytes). This makes it easier for users to assess and analyze CDN traffic based on report data. In order to display data as a larger unit, it must undergo a conversion process from Bytes to Kilobytes to Megabytes to Gigabytes to Terabytes. The conversion factor used in this calculation is 1000 instead of 1024. This has the benefit of simplifying the reporting interface and eliminating confusion for general users in differences between decimal and binary representation of these values.

Note: The above conversion process complies with the convention set by IEC in accordance with the addendum to IEC 60027-2, IEE, and ISO standards.

Note: Our reports do not use the kibi, mebi, gibi terminology. Instead, our reports use the correct terminology of KB, MB, GB, and TB.

Exporting Report Data


Report data, which is displayed directly below a graph, can be exported as a comma-separated values (CSV) file. This CSV file contains a list of headers and the report data that corresponds to them. This industry-standard file format is widely supported by log file analysis tools (e.g., Sawmill) and spreadsheet applications (e.g., Microsoft Excel). This allows you to analyze report data using your preferred software application.

Tip: If you plan on using a third-party application to analyze CDN activity data, then an alternative approach may be to use raw log files as a data source. Raw log files can be archived either on our CDN storage service or on a server of your choosing. For additional information on how to archive raw log data, please refer to the **MCC User Guide**.

Note: Report data cannot be exported for certain reports, such as the Traffic Summary report and reports generated for the Real-Time Statistics module.

Note: The statistical information used to generate our reports is mined from data logged by our servers.

To export report data

1. Generate the desired report.
2. Click the spreadsheet icon (). Report data will be saved on your computer as a CSV file.
3. Import the CSV file using the desired log file analysis or spreadsheet application.

Tip: By default, a generic report-specific name will be assigned to the CSV file. Therefore, it is recommended that you rename the CSV file immediately after downloading it.

Core Reporting Module

Overview

A variety of basic reports are available to help you view usage patterns and help you improve how our CDN serves assets to your clients. View a list of these reports by finding the **Analytics** menu and then selecting **Core Reports**. The side navigation bar will list reports that provide usage information on a variety of metrics, such as traffic, bandwidth, and CDN storage space. Cache statistics are also available, which allow you to check whether caching has been optimized to ensure quick data delivery speeds. Each report in the Core Reporting module is described in this chapter.

Note: This type of report can only be generated for CDN activity within the last 18 months.

Logged Data

Under certain circumstances, it may take up to 7 days to gather a comprehensive set of log data for all CDN activity associated with your account. Viewing a report under such circumstances would provide a partial representation of the activity that took place during the last 7 days. A typical cause for the backlogging of report data occurs when streams are left open for long periods of time or when a server is taken down for maintenance.

Platform-Specific Reports

The majority of the reports in the Core Reporting module allow you to generate a report by selecting a platform. The following table describes the available options and the platform and protocol(s) for which a report will be generated.

Report Name	Platform	Protocol(s)
HTTP Large	HTTP Large	HTTP
HTTPS Large	HTTP Large	HTTPS
HTTP Small	HTTP Small	HTTP
HTTPS Small	HTTP Small	HTTPS
ADN	Application Delivery Network	HTTP
ADN SSL	Application Delivery Network	HTTPS

Traffic Summary

The Traffic Summary report provides data on how much of your traffic is being routed through our CDN for each platform.

Note: This report reflects the CDN traffic used to bill your account.

Note: For the purposes of billing, report data is closed on the 3rd of the month. This means that report data for the current month is incomplete until after the third of the next month.

When viewing traffic data, you can choose to view all traffic or only the traffic that is routed to a particular region (e.g., North America & Europe or Asia & South America).

Note: Traffic routed through premium POPs in Asia, Australia, Latin America, and Nordic countries is billed at a different rate. The geographic delivery region associated with purchased services determines whether your traffic can leverage premium transit links and POPs.

The screenshot displays the Traffic Summary report interface. On the left is a navigation menu with options like Traffic Summary, Bandwidth, Data Transferred, Hits, Cache Statuses, Cache Hit Ratio, Cnames, CDN Storage, IPv4/IPv6, and Notes. The main content area shows 'Traffic: Global' and a help center link. Below this, there are filters for 'Regions' (set to Global) and 'Date Range' (set to October, 2014). The report includes two tables: 'Traffic By Media Type' and 'Customer Storage'.

Media Type Name	95% Bandwidth	Data Transferred (GB)	Data Transferred (TB)
HTTP Large Object	157679.0112 Mbps	15610653.29 GB	15610.65 TB
HTTPS Large Object	0.0001 Mbps	0.00 GB	0.00 TB
HTTP Small Object	7643.9163 Mbps	833072.43 GB	833.07 TB
HTTPS Small Object	239.9111 Mbps	22950.87 GB	22.95 TB
FMS	0.0000 Mbps	0.00 GB	0.00 TB
WMS	0.0000 Mbps	0.00 GB	0.00 TB
ADN	0.0000 Mbps	0.00 GB	0.00 TB
ADN SSL	0.0000 Mbps	0.00 GB	0.00 TB
Total	165562.8387 Mbps	16466676.58 GB	16466.68 TB

Current (MB)	Current (GB)	Current As Of	Max During Selected Month (MB)	Max During Selected Month (GB)
143235.07 MB	143.24 GB	2014-10-15	185409.27 MB	185.41 GB

Traffic Summary Report for the Global Region

The Traffic Summary report consists of two sections, which are Traffic By Media Type and Customer Storage. Both sections are explained below.

The Traffic By Media Type section provides a breakdown of total traffic activity by platform for the given region over the specified time period. A description is provided below for each field in the Traffic By Media Type portion of the report.

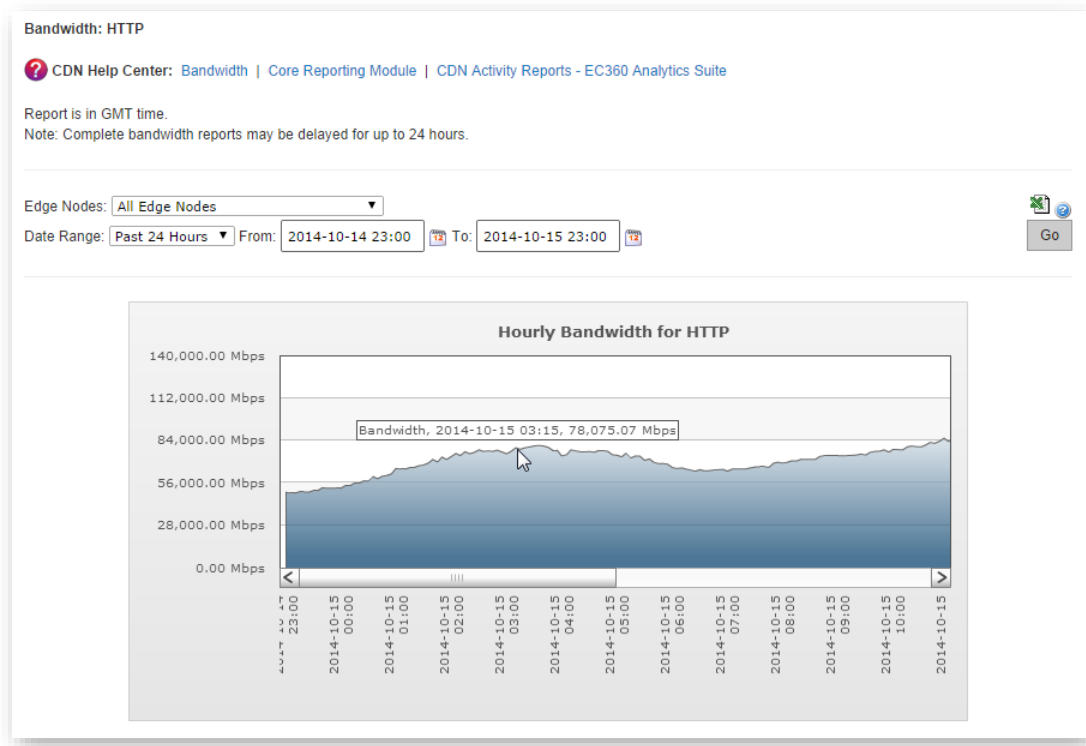
Field	Description
Media Type Name	Provides a list of platforms for which traffic information will be displayed.
95% Bandwidth	If your contract stipulates that you will be charged by bandwidth usage, then this field can be used as a billing indicator. It reports the amount of data (Megabits) transferred per second when your traffic is at 95% of peak usage.
Data Transferred (GB)	Indicates the total amount of data (GB) transferred through the CDN for the current customer on a particular platform.
Data Transferred (TB)	Indicates the total amount of data (TB) transferred through the CDN for the current customer on a particular platform.

The Customer Storage section provides CDN storage statistics. A description is provided below for each field in the Customer Storage portion of the report.

Field	Description
Current (MB)	Indicates the current amount of CDN storage space used in megabytes (MB). The region and time period selected for the Traffic Summary report do not affect this option.
Current (GB)	Indicates the current amount of CDN storage space used in gigabytes (GB). The region and time period selected for the Traffic Summary report do not affect this option.
Current As Of	Indicates the date (YYYY-MM-DD) used to calculate CDN storage statistics by the Current (MB) and the Current (GB) fields.
Max During Selected Month (MB)	Indicates the maximum amount of CDN storage spaced used in megabytes (MB) for the time period defined in the Date Range option.
Max During Selected Month (GB)	Indicates the maximum amount of CDN storage spaced used in gigabytes (GB) for the time period defined in the Date Range option.

Bandwidth

This type of report allows you to view bandwidth (Mbps) usage by platform and protocol. This report is based on traffic statistics tracked by our edge servers.



Bandwidth Report for the HTTP Large Platform

A Bandwidth report consists of a graph indicating the bandwidth usage for the selected platform, protocol, and/or POP(s) over a particular time period.

Tip: You can view the amount of bandwidth usage at any given time by hovering over the desired point in the line (as illustrated above). The time interval between points in the line is determined by whether you are viewing an hourly (5 minute increments) or daily (hourly increments) report.

Note: Data is reported in 5 minute chunks.

Note: If you are viewing data for a recent time period (e.g., Today or Past 24 Hours), then you will notice that amount of bandwidth usage tapers off as it approaches the current time. This trend is a result of the amount of time that it takes for log information to be collected from our edge servers.

The data that was used to generate the graph can be viewed below it. A table indicates bandwidth usage (megabits per second) in five minute intervals over the time period covered by the report.

Graph information:


- **Title:** Indicates whether time will be expressed by hour or by day.
- **Y-axis (Left):** Indicates bandwidth usage in megabits per second (Mbps).
- **X-axis (Below):** Indicates the time for the bandwidth reported by the corresponding data point in the graph. The time interval between data points can be either hourly or daily.
 - **Hourly:** This time interval is used for the "Past 24 Hours," "Today," and "Yesterday" time periods.
 - **Format:** YYYY-MM-DD hh:mm (e.g., 2015-12-12 12:00)
 - 24 hour format
 - UTC/GMT time zone
 - **Daily:** This time interval is used for all other time periods.
 - **Format:** YYYY-MM-DD (e.g., 2015-12-12)

To generate a Bandwidth report

1. Select the desired platform and protocol combination (e.g., HTTPS Large) from the side navigation bar. A graph will display bandwidth data for that platform over the last 24 hours for all POPs.
2. Optional. Filter the report to only show data for a particular POP (a.k.a., edge node).
3. Optional. Filter the report to show data for a different time period. You can choose to select a predefined time period or specify the desired date range.

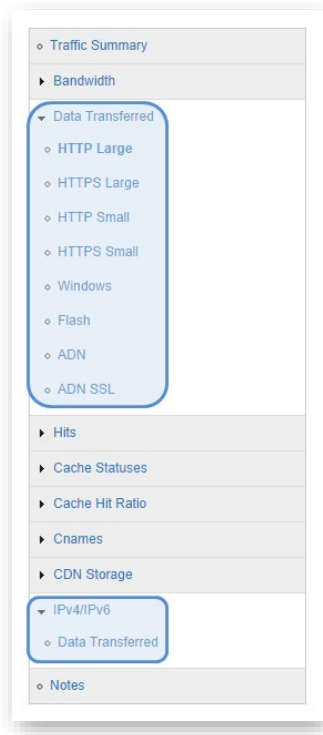
Data Transferred

This type of report allows you to view the amount of data transferred (MB) on your account. This report is based on traffic statistics tracked by our edge servers.



Data Transferred Report for the HTTP Large Platform

A Data Transferred report can report data based on platform/protocol or by the type of IP address (i.e., IPv4 or IPv6) used by clients to request your content. The manner in which data is organized for this type of report is determined by your selection in the side navigation bar.



Core Reports Side Navigation Bar (Outlined Areas Indicate Data Transferred Reports)

A Data Transferred report consists of a graph indicating the amount of data transferred for the selected platform/protocol or IP protocol over a particular time period.

Tip: You can view the amount of data transferred at a specified point in time by hovering over the desired bar in the graph. The time interval represented by each bar is determined by whether you are viewing an hourly (5 minute increments) or daily (hourly increments) report.

Note: If you are viewing report data by platform, then you can also customize it by filtering the report to only show data for a particular POP (a.k.a., edge node).

Note: This operation excludes data for transactions that did not complete during the requested time period, even if the transaction started before or during the time period covered by the report.

Note: Data is reported in 5 minute chunks.

Note: If you are viewing data for a recent time period (e.g., Today or Past 24 Hours), then you will notice that amount of data transferred tapers off as it approaches the current time. This trend is a result of the amount of time that it takes for log information to be collected from our edge servers.

The data that was used to generate the graph can be viewed below it. A table indicates the amount of data transferred in megabytes and gigabytes in five minute intervals over the time period covered by the report.

Graph information:

- **Title:** Indicates whether time will be expressed by hour or by day.
- **Y-axis (Left):** Indicates the amount of data transferred in megabytes (MB).
- **X-axis (Below):** Indicates the time for the data transferred reported by the corresponding bar in the graph. The time interval between bars can be either hourly or daily.
 - **Hourly:** Information on when this time interval is used and its format is provided below.
 - **Scope:** Time interval varies according to report type.
 - **Platform/Protocol:** This time interval is used for the "Past 24 Hours," "Today," and "Yesterday" time periods.
 - **IP Protocol:** This time interval is used for all time periods.
 - **Format:** YYYY-MM-DD hh:mm (e.g., 2015-12-12 12:00)
 - 24 hour format
 - UTC/GMT time zone
 - **Daily:** This time interval is used for all other time periods for reports generated based on platform/protocol.
 - **Format:** YYYY-MM-DD (e.g., 2015-12-12)

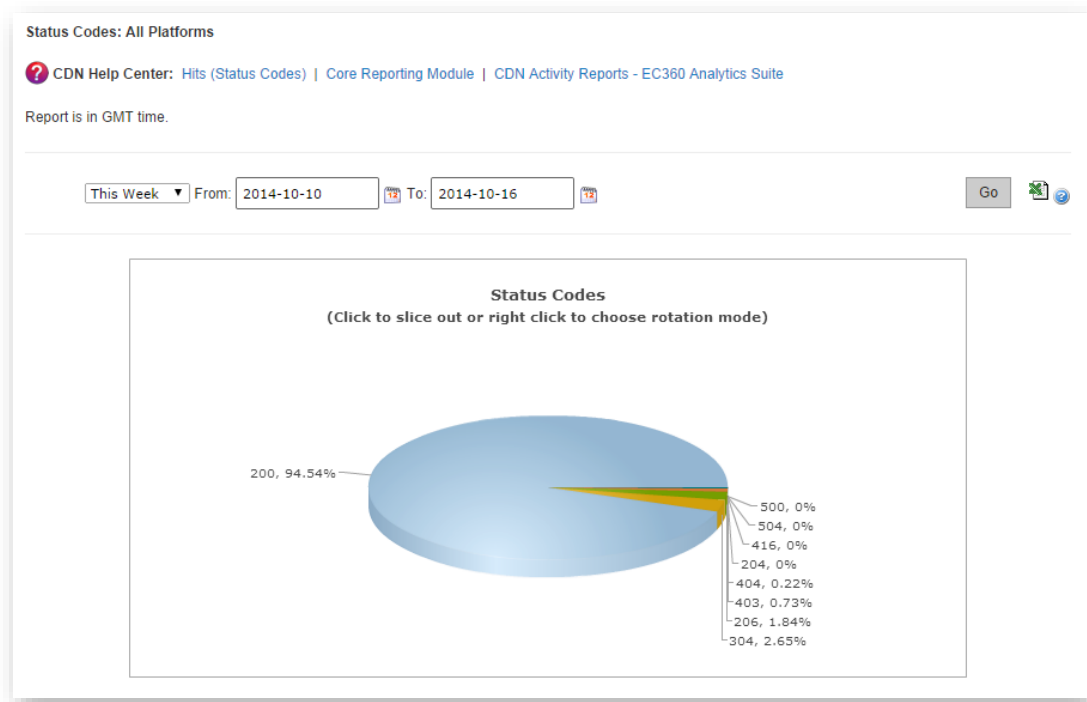
To generate a Data Transferred report

1. Generate a report based organized according to:
 - **Platform/Protocol:** Select the desired platform and protocol combination (e.g., HTTPS Large) from the side navigation bar. A graph will display data transferred statistics for that platform over the last 24 hours for all POPs.
 - **IP Protocol:** Select the desired IP protocol (i.e., IPv4 or IPv6) from the side navigation bar. A graph will display data transferred statistics for that IP protocol over the last 24 hours for all POPs.
2. Optional. Use the **Edge Nodes** option to filter platform-specific reports to only show data for a particular POP (a.k.a., edge node).
3. Optional. Filter the report to show data for a different time period. You can choose to select a predefined time period or specify the desired date range.

Hits (Status Codes)

Every request for content will generate an HTTP status code. This status code describes how an edge server, origin shield server, or origin server handled the request. For example, 2xx status codes indicate that the request was successfully served to a client, while a 4xx status code indicates that an error took place.

The Hits report allows you to view the frequency of each status code as a percentage of total requests. This report is based on traffic statistics tracked by our edge servers.



Hits Report for the HTTP Large Platform

The data that was used to generate the pie chart can be viewed directly below it. The following information will be reported for each status code: a description of the status code, the total number of hits generated by that status code, and the percentage of total hits for which it was returned.

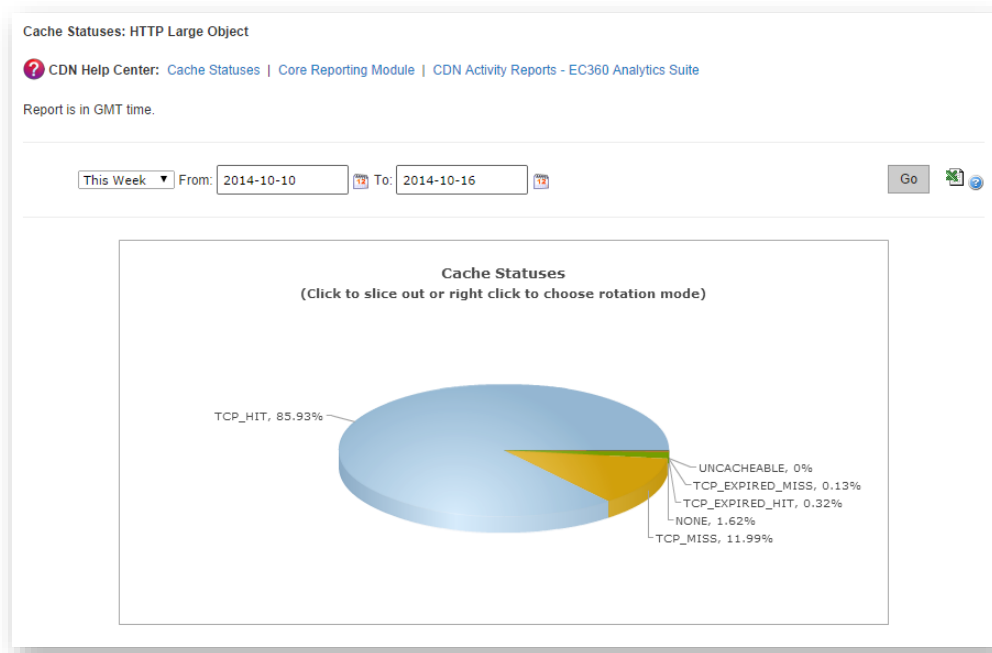
To generate a Hits (Status Codes) report

1. Select "All Platforms" or the desired platform from the side navigation bar. A graph will display status code statistics for the selected platform(s) over the last week for all POPs.
2. Optional. Filter the report to show data for a different time period. You can choose to select a predefined time period or specify the desired date range.

Cache Statuses

The Cache Statuses report gives a detailed breakdown of cache behavior, which may reveal approaches for improving the overall end-user experience. Since the fastest performance comes from cache hits, you can optimize data delivery speeds by minimizing cache misses and expired cache hits. Cache misses can be reduced by configuring your origin server to avoid assigning "no-cache" response headers, by avoiding query-string caching except where strictly needed, and by avoiding uncacheable response codes. Expired cache hits can be avoided by making an asset's max-age as long as possible to minimize the number of requests to the origin server.

Note: For a detailed explanation of each cache status, please refer to **Appendix C: Cache Statuses**.



Cache Statuses Report for All Platforms

The data that was used to generate the pie chart can be viewed directly below it. Each cache status is listed along with a description, the total number of times that it occurred, and the percentage of requests that resulted in this cache status.

To generate a Cache Statuses report

1. Select "All Platforms" or the desired HTTP-based platform from the side navigation bar. A graph will display cache status statistics for the selected platform(s) over the last week for all POPs.
2. Optional. Filter the report to show data for a different time period. You can choose to select a predefined time period or specify the desired date range.

Cache Hit Ratio

The purpose of the Cache Hit Ratio report is to indicate the percentage of cacheable requests that were served directly from cache to the requester. In other words, the percentage of requests that met the following requirements:

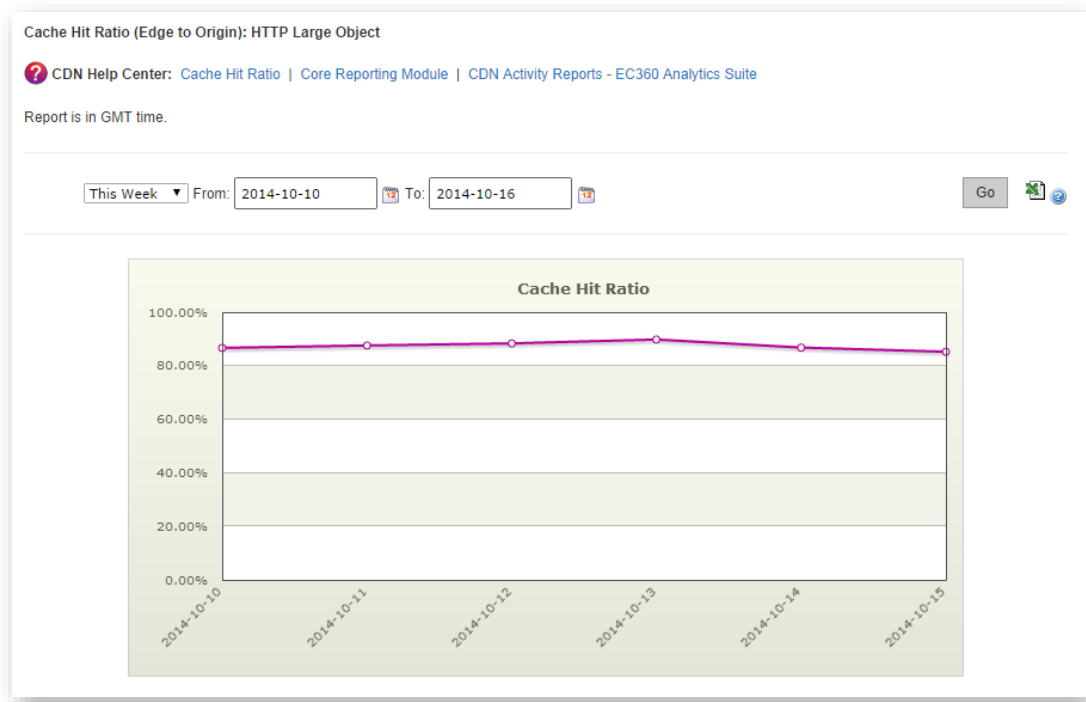
- The requested content was cached on the POP closest to the requester.
- The request was served directly from the edge of our network.
- The request did not require revalidation with the origin server.

This report excludes all of the following cases:

- Requests that are denied due to HTTP Rules Engine.
- Requests that are denied due to Token-Based Authentication.
- Requests that are denied due to country filtering options.
- Requests that contain query string URLs when the query-string caching feature has been set to "no-cache."
- Requests for assets whose headers indicate that they should not be cached. For example, Cache-Control: private, Cache-Control: no-cache, or Pragma: no-cache headers will prevent an asset from being cached.
 - An exception to the above scenario may occur when HTTP Rules Engine is used to override the requested content's cache policy.
- Requests for partially cached content.

This report consists of a graph and statistical information on cache hits. The formula through which cache hit percentage is calculated is described below.

$$\text{Cache Hit Percentage} = (\text{TCP_HIT} / (\text{TCP_HIT} + \text{TCP_MISS})) * 100$$



Cache Hit Ratio Report for All Platforms

Tip: View the exact cache hit ratio on a given date by hovering over the hollow circles in the line graph.

The data that was used to generate the graph can be viewed directly below it. The date, total number of hits, total number of misses, and the cache hit percentage will be displayed for each plotted point in the graph.

Graph information:

- **Y-axis (Left):** Measures cache hits as a percentage of overall requests.
- **X-axis (Below):** Indicates the time for the cache hit ratio reported by the corresponding data point in the graph. The time interval between data points is always daily.
 - **Format:** YYYY-MM-DD (e.g., 2015-12-12)

To generate a Cache Hit Ratio report

1. Select "All Platforms" or the desired HTTP-based platform from the side navigation bar. A line graph will illustrate the percentage of hits that result in an asset being served from cache for the selected platform(s) over the last week for all POPs.
2. Optional. Filter the report to show data for a different time period. You can choose to select a predefined time period or specify the desired date range.

Cnames (Deprecated)

Important: All five types of Cnames reports (i.e., All Platforms, HTTP Large, HTTP Small, Flash, and ADN) have been deprecated as of 10/31/2015. Data is no longer collected for these reports and support for them will slowly be phased out. However, for the purpose of viewing historical data, these reports will remain available for a reasonable time period after end-of-life.


Tip: The Custom Reports module may be used to generate reports on edge CNAMEs. A prerequisite for this report is the identification of the edge CNAME configurations for which report data will be logged. For additional information, please refer to the **Custom Reports** chapter below.

The Cnames report provides data transferred statistics for each edge CNAME.

CDN Storage Usage

The CDN Storage Usage report provides daily statistics on CDN storage usage.

Note: CDN storage usage information is collected once a day. As a result, it will not reflect fluctuations in CDN storage usage within a given day.



CDN Storage Report

The data that was used to generate the graph can be viewed below the graph. A date entry and the amount of data that was stored at that time will be displayed for each data point in the graph.

Graph information:

- **Y-axis (Left):** Indicates how much data was stored on a CDN origin server.
- **X-axis (Below):** Indicates the time at which data storage on a CDN origin server was at a certain level. The time interval between data points is always daily.
 - **Format:** YYYY-MM-DD (e.g., 2015-12-12)

To generate a CDN Storage Usage report

1. Expand **CDN Storage** from the side navigation bar and then click **Usage**.
2. Optional. Filter the report to show data for a different time period. You can choose to select a predefined time period or specify the desired date range.

Edge Image Optimizer Usage

The Edge Image Optimizer Usage report tracks the number of requests that required Edge Image Optimizer processing over the requested time period.

Note: This report excludes requests for previously processed images that were served from cache.

View this report's raw data directly below the chart.

To generate an Edge Image Optimizer Usage report

1. Load the Edge Image Optimizer report.
 - i. From the main menu, navigate to **Analytics | Core Reports**.
 - ii. From the side navigation bar, expand **Edgelets** and then select **Edge Image Optimizer Usage**.

By default, this report loads a graph that tracks the number of requests that required image processing over the last week.

2. Optional. Use the **Date Range** option to define either a custom or predefined (e.g., This Week) time period for which report data will be generated.

DNS

Note: The DNS section of our Core Reporting module contains report data for our Route solution. It does not report DNS statistics for CDN data delivery.

The Route Summary Query Count report indicates the total number of DNS queries that were resolved by our Route name servers for each of your fully managed and secondary zones. It also indicates the total number of DNS requests that were resolved for all of your zones.

To generate a Route Summary Query Count report

1. Expand "DNS" from the side navigation bar and then select "Route Summary Query."
2. Optional. Filter the report to show data for a different time period. You can choose to select a predefined time period or specify the desired date range.

Custom Reports

Overview

Unlike other reporting modules, the Custom Reports module allows you to define the type of requests for which data will be collected. This process requires the activation of the Custom Reports capability, which allows our edge servers to keep track of the number of hits and the amount of data transferred for all requests that match the specified request criteria. After which, generate a report on this information by finding the **Analytics** menu and then selecting **Custom Reports**.

Note: The primary function of these reports is to assess performance. They should not be used for billing purposes or exact numeric statistics.

Note: This type of report can only be generated for CDN activity within the last 18 months.

Logging

The custom report capability must be activated on the desired types of requests before our servers will log the number of hits and the amount of data transferred for the requested criteria (e.g., edge CNAME). This differentiates it from other reports in the Analytics Suite. As a result, it is important to consider the following notes:

- For the purpose of generating custom reports, our servers will only log CDN usage data while the custom report capability is active for the desired types of requests. This means that data will not be logged for those requests prior to the activation date or after it has been deactivated.
- A custom report can be generated for a time period that includes one or more dates on which logging for the desired type of request was not active. This will generate a report that contains partial or no usage data.
- Custom report-specific logging for a particular type of request can be disabled at any time. This will not affect your ability to view historical data for the time period during which logging for that request type was turned on.
- The Custom Reports module tracks hits and data transferred by cache status code and HTTP status code. This capability was introduced on 8/04/2015. Data collected before that date does not track cache status and HTTP status codes. Therefore, older data will only be reported in the Hits, Data Transferred, and the Unassigned fields.

Important: The logging described above only affects the data that will be used to generate custom reports. Keep in mind that CDN usage statistics can always be viewed from one of our other reporting modules (e.g., Core Reporting).

Edge CNAMEs

This type of report provides hits and data transferred statistics for edge CNAMEs on which custom report logging has been enabled.

Tip: Custom report logging may be enabled on the desired edge CNAME from the **Edge CNAMEs** page.

Key information about activating custom reports on a per edge CNAME basis:

- Custom report data logging starts one hour after enabling an edge CNAME's custom reporting capability.
- Report data may be viewed by generating an Edge CNAMEs report for all platforms or a specific one. The coverage for this report will be limited to the edge CNAMEs for which custom report data was collected during the specified time period.

Report Details

Generate a custom report by defining both of the following:

- **Platform:** Choose whether a custom report will be generated for all platforms or a specific one. By default, the **Custom Reports** page generates a report for all platforms. Generate a platform-specific report by selecting the desired platform from the side navigation bar.
- **Type:** A custom report for hits or data transferred may be generated. By default, a custom report that tracks hits over a given time period (e.g., this week) will be generated. Generate a data transferred report by selecting "Data Transferred" from the **Metrics** option.
- **Statistics:** The type of statistics that will be shown below the bar chart is determined by the **Groupings** option. By default, a custom report will break down statistics by HTTP status codes.
- **Time Period:** The time period for which a custom report will be generated may be defined.

After generating a custom report, a bar chart will be generated for the top 10 edge CNAMEs according to the metric defined in the **Metrics** option.

The following variables may affect this calculation:

- A custom report will only include edge CNAMEs for which this capability has been activated.
- Was the custom report capability turned on for a particular edge CNAME during or after the time period specified for the custom report? If so, then the edge CNAME may either be ranked lower than it should or it may not even be included in the custom report.

This bar chart allows you to quickly assess which edge CNAMEs produce the most amount of traffic. The left-hand side of the graph (y-axis) indicates how much data was transferred per edge CNAME. Directly below the graph (x-axis), you will find a label for each of the top 10 edge CNAMEs.

Tip: View the amount of data transferred on a specific edge CNAME over the specified time frame by hovering over the desired bar.

Directly below the bar chart, statistics on hits or the amount of data transferred on a per edge CNAME basis are broken down by either cache status code or HTTP status code. The type of statistics that will be shown is determined by the **Groupings** option. A description is provided for each of these metrics in the **Appendix A: Custom Report Fields**.

Advanced Content Analytics

Overview

Similar to the Core Reporting module, the Advanced Content Analytics module contains reports that provide statistical information on CDN activity. However, the reports provided in the Advanced Content Analytics module allow you to probe a little bit deeper into where your content is being requested and what happens after it gets requested. Thus, helping you identify and analyze CDN usage patterns.

Reports & Log Collection

CDN activity data must be collected by the Advanced Content Analytics module before it can generate reports on it. This collection process occurs once a day and it covers the activity that took place during the previous day. This means that a report's statistics represent a sample of the day's statistics at the time it was processed, and do not necessarily contain the complete set of data for the current day. The primary function of these reports is to assess performance. They should not be used for billing purposes or exact numeric statistics.

Tip: If billing is your primary purpose for viewing a report, then you should generate and analyze reports (e.g., Traffic Summary, Bandwidth, or Data Transferred) in the Core Reporting module.

Note: The raw data from which Advanced Content Analytic reports are generated is available for at least 90 days.

Asset/Directory Location

Certain reports (e.g., By File, By File Detail, By Directory, etc.) need to specify a unique path to an asset or a directory. They are able to do so through the use of a CDN URL path. A CDN URL path starts with the content access point. This involves truncating the protocol and hostname from the CDN URL. The CDN URL path for the following sample CDN URL is indicated below.

Note: For the purposes of the Advanced Content Analytics module, edge CNAME URLs are converted to their equivalent CDN URLs. This allows an accurate tally for all statistics associated with an asset regardless of the CDN or edge CNAME URL used to request it.

Sample CDN URL:

`http://wac.0001.edgecastcdn.net/000001/main/index.html`

Sample CDN URL path:

/000001/main/index.html

The above sample CDN URL path starts with a content access point (e.g., /000001). A content access point (i.e., *yyxxxx* or *80xxxx/CustomerOrigin*) starts with an origin identifier. In this particular case, the first two digits of an origin identifier indicate whether the asset is stored on a CDN (i.e., 00) or a customer origin server (i.e., 80).

Advanced HTTP and ADN Reports

This section explains each report in the Advanced Content Analytics module. These reports provide detailed information on CDN activity for the HTTP Large and ADN. These reports may be accessed via the following menu items:

- **Advanced HTTP Reports:** This menu item contains the set of reports that provides graphs and statistics for the HTTP Large platform.
- **Advanced ADN Reports:** This menu item contains the set of reports that provides graphs and statistics for the ADN platform.

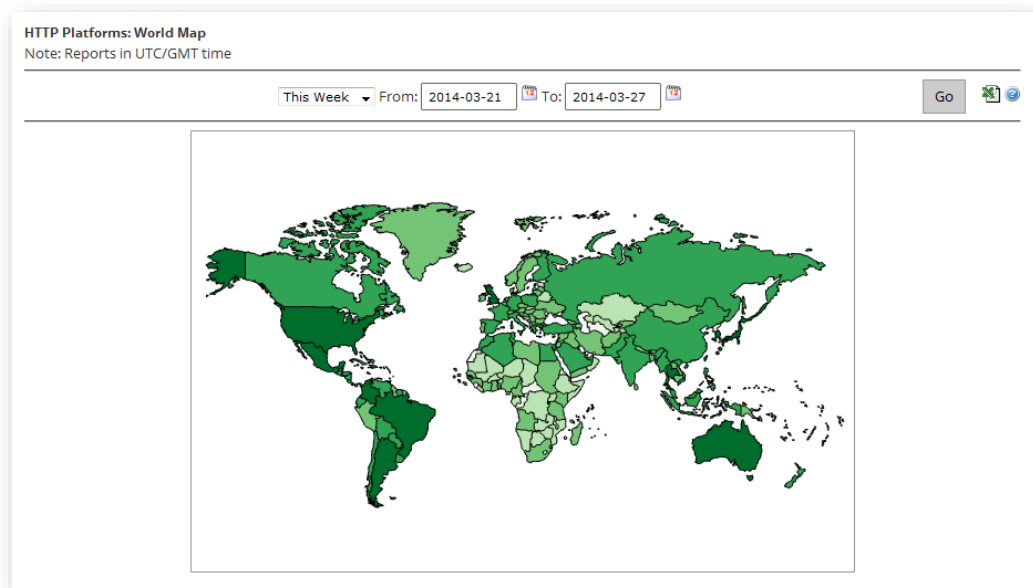
Note: Although the data contained in each of these reports are platform-specific, they are quite similar in substance and appearance. All platform-specific differences will be indicated when appropriate.

Geography Reports (Map-Based)

There are five reports that take advantage of a map to indicate the regions from which your content is being requested. These reports are World Map, United States Map, Canada Map, Europe Map, and Asia Pacific Map.

Note: The location from which a request originated is determined by looking up the client's IP address in a GeoIP database.

Each map-based report ranks geographic entities (i.e., countries, states, and provinces) according to the percentage of hits that originated from that region. Additionally, a map is provided to help you visualize the locations from which your content is being requested. It is able to do so by color-coding each region according to the amount of demand experienced in that region. Lighter shaded regions indicate lower demand for your content, while darker regions indicate higher levels of demand for your content.



Map-Based Report (World Map)

Detailed traffic and bandwidth information for each region is provided directly below the map. This allows you to view the total number of hits, the percentage of hits, the total amount of data transferred (in gigabytes), and the percentage of data transferred for each region. A description is provided for each of these metrics in the **Appendix A: Advanced Content Analytics Fields**. Finally, when you hover over a region (i.e., country, state, or province), the name and the percentage of hits that occurred in the region will be displayed as a tooltip.


A brief description is provided below for each type of map-based geography report.

Report Name	Description
World Map	This report allows you to view the worldwide demand for your CDN content. Each country is color-coded on the world map to indicate the percentage of hits that originated from that region.
United States Map	This report allows you to view the demand for your CDN content in the United States. Each state is color-coded on this map to indicate the percentage of hits that originated from that region.
Canada Map	This report allows you to view the demand for your CDN content in Canada. Each province is color-coded on this map to indicate the percentage of hits that originated from that region.
Europe Map	This report allows you to view the demand for your CDN content in Europe. Each country is color-coded on this map to indicate the percentage of hits that originated from that region.
Asia Pacific Map	This report allows you to view the demand for your CDN content in Asia. Each country is color-coded on this map to indicate the percentage of hits that originated from that region.

Geography Reports (Bar Charts)

There are two additional reports that provide statistical information according to geography, which are Top Cities and Top Countries. These reports rank cities and countries, respectively, according to the number of hits that originated from those regions. Upon generating this type of report, a bar chart will indicate the top 10 cities or countries that requested content over the HTTP Large or ADN platform. This bar chart allows you to quickly assess the regions that generate the highest number of requests for your content.

Reminder: The location from which a request originated is determined by looking up the client's IP address in a GeolIP database.



Top Cities Report

The left-hand side of the graph (y-axis) indicates how many hits occurred in the specified region. Directly below the graph (x-axis), you will find a label for each of the top 10 regions.

Tip: If you hover over a bar, the name and the total number of hits that occurred in the region will be displayed as a tooltip.

Note: The tooltip for the Top Cities report identifies a city by its name, state/province, and country abbreviation.

Note: If the city or region (i.e., state/province) from which a request originated from could not be determined, then it will indicate that they are unknown. If the country is unknown, then two question marks (i.e., ??) will be displayed.

Note: A report may include metrics for "Europe" or the "Asia/Pacific Region." Those items are not meant to provide statistical information on all IP addresses in those regions. Rather, they only apply to requests that originate from IP addresses that are spread out over Europe or Asia/Pacific instead of to a specific city or country.

The data that was used to generate the bar chart can be viewed below it. There you will find the total number of hits, the percentage of hits, the amount of data transferred (in gigabytes), and the percentage of data transferred for the top 250 regions. A description is provided for each of these metrics in the **Appendix A: Advanced Content Analytics Fields**.

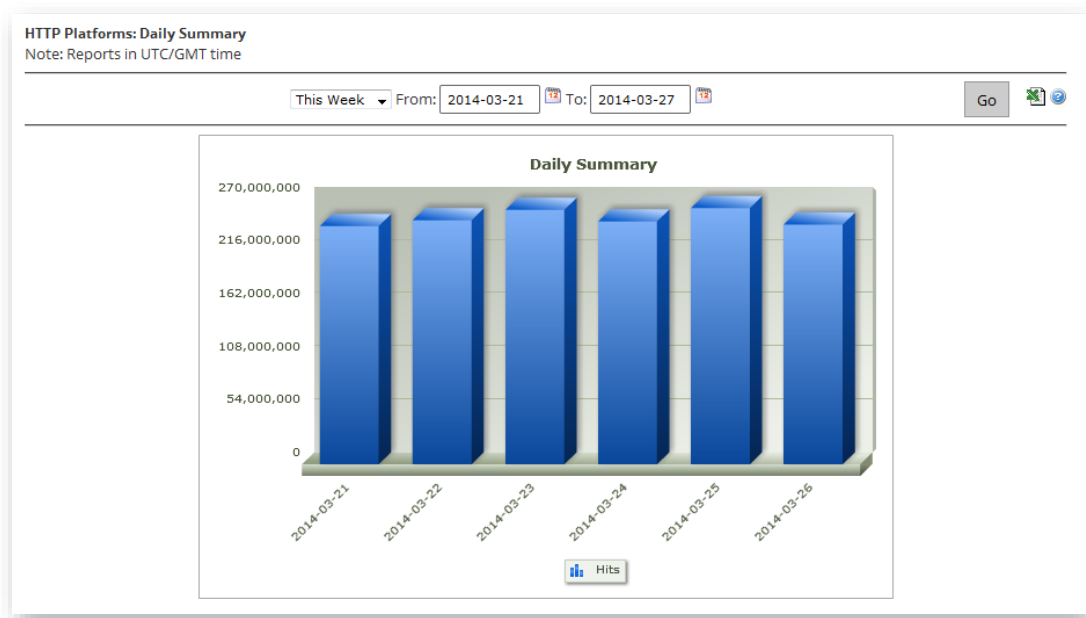
A brief description is provided for both types of reports below.

Report Name	Description
Top Cities	This report ranks cities according to the number of hits that originated from that region.
Top Countries	This report ranks countries according to the number of hits that originated from that region.

Daily Summary Report

The Daily Summary report allows you to view the total number of hits and data transferred over the HTTP Large or ADN platform on a daily basis. This information can be used to quickly discern CDN activity patterns. For example, this report can help you detect which days experienced higher or lower than expected traffic.

Upon generating this type of report, a bar chart will provide a visual indication as to the amount of platform-specific demand experienced on a daily basis over the time period covered by the report. It will do so by displaying a bar for each day in the report. For example, selecting the time period called "Last Week" will generate a bar chart with seven bars. Each bar will indicate the total number of hits experienced on that day.



Daily Summary Report

The left-hand side of the graph (y-axis) indicates how many hits occurred on the specified date. Directly below the graph (x-axis), you will find a label that indicates the date (YYYY-MM-DD) for each day included in the report.

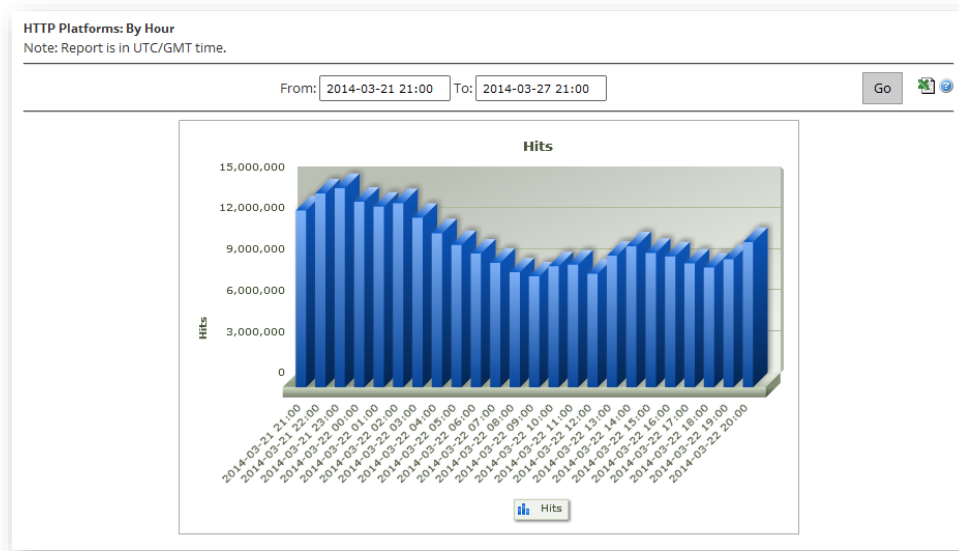
Tip: If you hover over a bar, the total number of hits that occurred on that date will be displayed as a tooltip.

The data that was used to generate the bar chart can be viewed below it. There you will find the total number of hits and the amount of data transferred (in gigabytes) for each day covered by the report. A description is provided for each of these metrics in the **Appendix A: Advanced Content Analytics Fields**.

By Hour Report

The By Hour report allows you to view the total number of hits and data transferred over the HTTP Large platform on an hourly basis. This information can be used to quickly discern CDN activity patterns. For example, this report can help you detect the time periods during the day that experience higher or lower than expected traffic.

Upon generating this type of report, a bar chart will provide a visual indication as to the amount of platform-specific demand experienced on an hourly basis over the time period covered by the report. It will do so by displaying a bar for each hour covered by the report. For example, selecting a 24 hour time period will generate a bar chart with twenty four bars. Each bar will indicate the total number of hits experienced during that hour.



By Hour Report

The left-hand side of the graph (y-axis) indicates how many hits occurred on the specified hour. Directly below the graph (x-axis), you will find a label that indicates the date/time (YYYY-MM-DD hh:mm) for each hour included in the report. Time is reported using 24 hour format and it is specified using the UTC/GMT time zone.

Tip: If you hover over a bar, the total number of hits that occurred during that hour will be displayed as a tooltip.


Note: Data is reported in 1 hour chunks.

The data that was used to generate the bar chart can be viewed below it. There you will find the total number of hits and the amount of data transferred (in gigabytes) for each hour covered by the report. A description is provided for each of these metrics in the **Appendix A: Advanced Content Analytics Fields**.

By File Report

The By File report allows you to view the amount of demand and the traffic incurred over the HTTP Large or ADN platform for the most requested assets. Upon generating this type of report, a bar chart will be generated on the top 10 most requested assets over the specified time period.

Reminder: For the purposes of this report, edge CNAME URLs are converted to their equivalent CDN URLs. This allows an accurate tally for the total number of hits associated with an asset regardless of the CDN or edge CNAME URL used to request it.



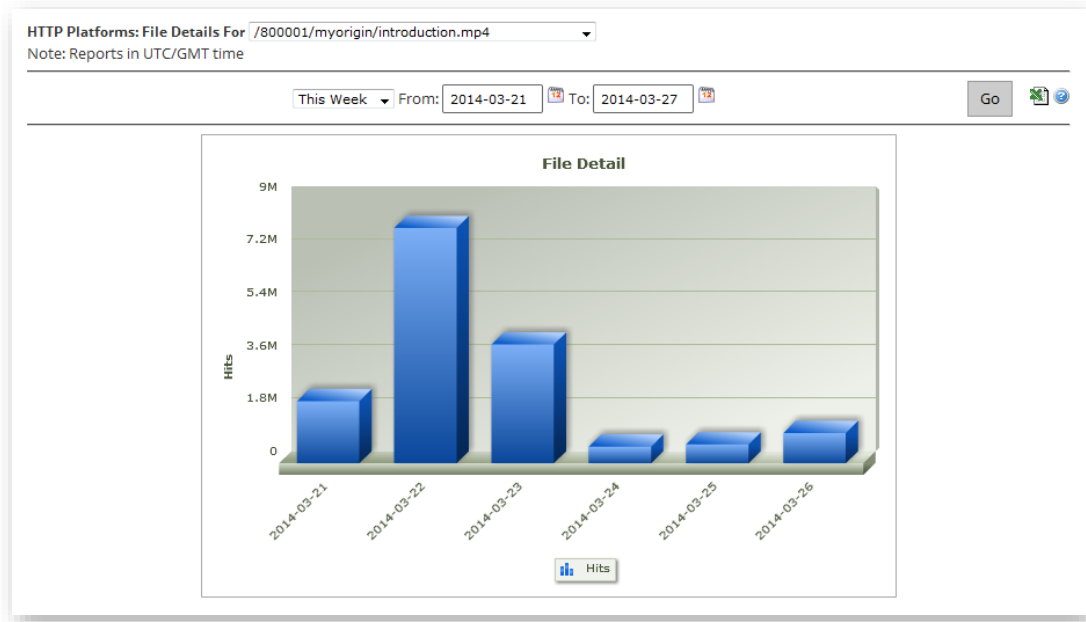
By File Report

The left-hand side of the graph (y-axis) indicates the number of requests for each asset over the specified time period. Directly below the graph (x-axis), you will find a label that indicates the file name for each of the top 10 requested assets.

The data that was used to generate the bar chart can be viewed below it. There you will find the following information for each of the top 250 requested assets: relative path, the total number of hits, the percentage of hits, the amount of data transferred (in gigabytes), and the percentage of data transferred. A description is provided for each of these metrics in the **Appendix A: Advanced Content Analytics Fields**.

By File Detail Report

The By File Detail report allows you to view the amount of demand and the traffic incurred over the HTTP Large or ADN platform for a particular asset. At the very top of this report is the **File Details For** option. This option provides a list of your most requested assets on the selected platform. In order to generate a By File Detail report, you will need to select the desired asset from the **File Details For** option. After which, a bar chart will indicate the amount of daily demand that it generated over the specified time period.



By File Detail Report

The left-hand side of the graph (y-axis) indicates the total number of requests that an asset experienced on a particular day. Directly below the graph (x-axis), you will find a label that indicates the date (YYYY-MM-DD) for which CDN demand for the asset was reported.


The data that was used to generate the bar chart can be viewed below it. There you will find the total number of hits and the amount of data transferred (in gigabytes) for each day covered by the report. A description is provided for each of these metrics in the **Appendix A: Advanced Content Analytics Fields**.

By File Type Report

The By File Type report allows you to view the amount of demand and the traffic incurred over the HTTP Large or the ADN platform by file type. Upon generating this type of report, a donut chart will indicate the percentage of hits generated by the top 10 file types.

Tip: If you hover over a slice in the donut chart, the Internet media type of that file type will be displayed as a tooltip.

Note: This report is available for the HTTP Large and ADN platforms.



By File Type Report for the HTTP Large platform

The data that was used to generate the donut chart can be viewed below it. There you will find the file name extension/Internet media type, the total number of hits, the percentage of hits, the amount of data transferred (in gigabytes), and the percentage of data transferred for each of the top 250 file types. A description is provided for each of these metrics in the **Appendix A: Advanced Content Analytics Fields**.

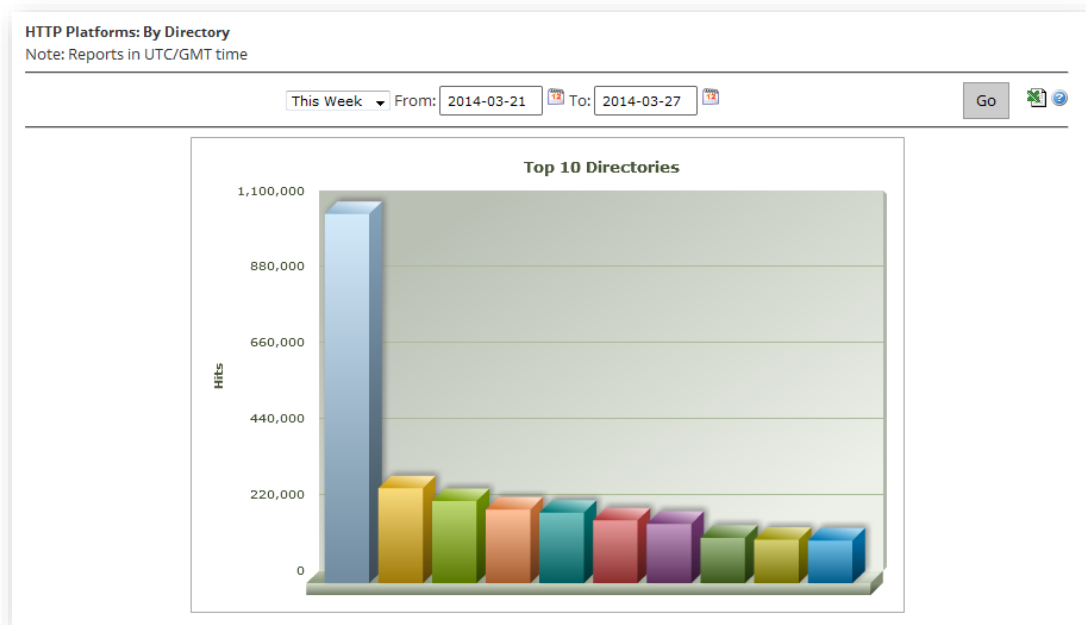
By Directory Report

The By Directory report allows you to view the amount of demand and the traffic incurred over the HTTP Large or ADN platform for content from a particular directory. Upon generating this type of report, a bar chart will indicate the total number of hits generated by content in the top 10 directories.

Tip: Hover over a bar to view the relative path to the corresponding directory.

Note: Content stored in a subfolder of a directory does not count when calculating demand by directory. This calculation relies solely on the number of requests generated for content stored in the actual directory.

Reminder: For the purposes of this report, edge CNAME URLs are converted to their equivalent CDN URLs. This allows an accurate tally for all statistics associated with an asset regardless of the CDN or edge CNAME URL used to request it.



By Directory Report

The left-hand side of the graph (y-axis) indicates the total number of requests for the content stored in your top 10 directories. Each bar on the chart represents a directory. Use the color-coding scheme to match up a bar to a directory listed in the **Top 250 Full Directories** section.

The data that was used to generate the bar chart can be viewed below it. There you will find the following information for each of the top 250 directories: relative path, the total number of hits, the percentage of hits, the amount of data transferred (in gigabytes), and the percentage of data transferred. A description is provided for each of these metrics in the **Appendix A: Advanced Content Analytics Fields**.

By Browser Report


The By Browser report allows you to view which browsers are used to request content from the HTTP Large platform. Upon generating this type of report, a pie chart will indicate the percentage of requests handled by the top 10 browsers.

Tip: Hover over a slice in the pie chart to view a browser's name and version.

Note: This report is available for the HTTP Large and the ADN platforms.

Note: For the purposes of this report, each unique browser/version combination is considered a different browser.

Note: The slice called "Other" indicates the percentage of requests handled by all other browsers and versions.

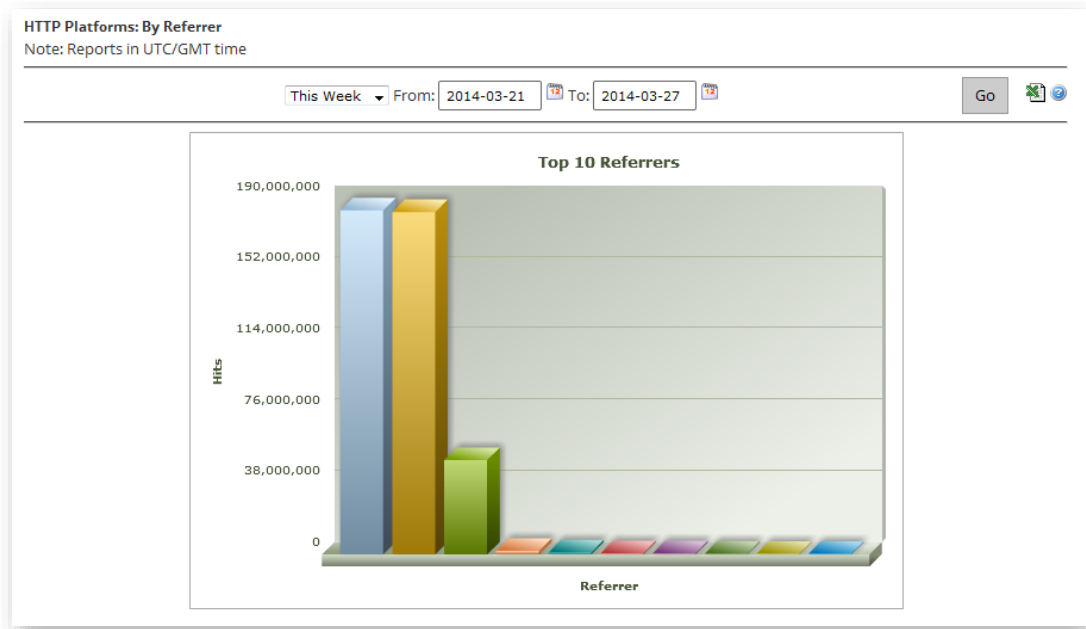


By Browser Report for the HTTP Large platform

The data that was used to generate the pie chart can be viewed below it. There you will find the browser type/version number, the total number of hits and the percentage of hits for each of the top 250 browsers. A description is provided for each of these metrics in the **Appendix A: Advanced Content Analytics Fields**.

By Referrer Report

The By Referrer report allows you to view the top referrers to content on the selected platform (i.e., HTTP Large or ADN). A referrer indicates the hostname from which a request was generated. Upon generating this type of report, a bar chart will indicate the amount of demand (i.e., hits) generated by the top 10 referrers.



By Referrer Report

The left-hand side of the graph (y-axis) indicates the total number of requests that an asset experienced for each referrer. Each bar on the chart represents a referrer. Use the color-coding scheme to match up a bar to a referrer listed in the **Top 250 Referrer** section.

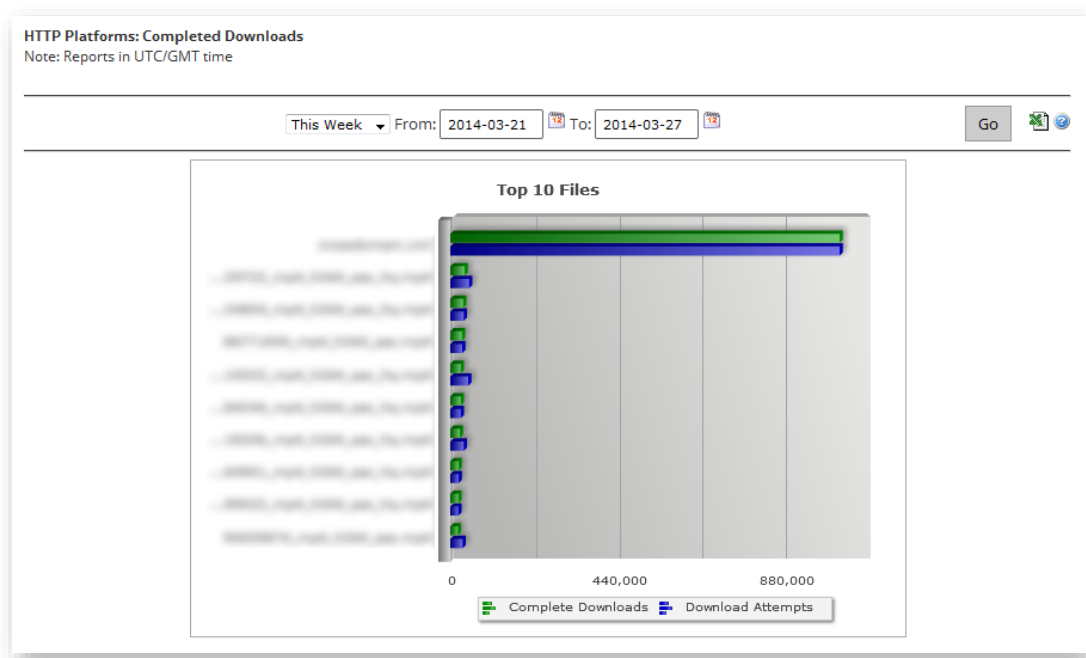
The data that was used to generate the bar chart can be viewed below it. There you will find the URL, the total number of hits, and the percentage of hits generated from each of the top 250 referrers. A description is provided for each of these metrics in the **Appendix A: Advanced Content Analytics Fields**.

By Download Report

The By Download report allows you to analyze download patterns for your most requested content. The top of the report contains a bar chart that compares attempted downloads with completed downloads for the top 10 requested assets. Each bar is color-coded according to whether it is an attempted download (blue) or a completed download (green).

Note: This report is available for the HTTP Large and the ADN platforms.

Reminder: For the purposes of this report, edge CNAME URLs are converted to their equivalent CDN URLs. This allows an accurate tally for all statistics associated with an asset regardless of the CDN or edge CNAME URL used to request it.



By Download Report for the HTTP Large platform

The left-hand side of the graph (y-axis) indicates the file name for each of the top 10 requested assets. Directly below the graph (x-axis), you will find labels that indicate the total number of attempted/completed downloads.

Directly below the bar chart, the following information will be listed for the top 250 requested assets: relative path (including file name), the number of times that it was downloaded to completion, the number of times that it was requested, and the percentage of requests that resulted in a complete download. A description is provided for each of these metrics in the **Appendix A: Advanced Content Analytics Fields**.

Calculating Attempted Downloads

Only the following types of requests are counted when calculating attempted downloads:

- Requests that result in a 200 OK.
- Byte-range requests for the start of a file (e.g., bytes=0-499) that result in a 206 Partial Content.

Calculating Completed Downloads

An HTTP client (e.g., web browser) does not inform our CDN service when an asset has been completely downloaded. As a result, we have to assess whether an asset has been completely downloaded according to status codes and byte-range requests.

Only requests that satisfy all of the following criteria count towards completed downloads:

Criterion	Description
Status Code	All requests that result in a 200 OK will count towards a completed download. Byte-Range Requests Only The request must result in a 206 Partial Content status code.
Byte-Range Requests Only Coverage	The set of byte-range requests (e.g., bytes=0-499, bytes=500-999, bytes=1000-1499, etc.) from the HTTP client must cover the entire length of the requested asset.
Byte-Range Requests Only File Size	The total amount of data transferred to the HTTP client must be equal to or greater than the asset's file size.

Due to the interpretive nature of this report, you should keep in mind the following points that may alter the consistency and accuracy of this report.

- Traffic patterns cannot be accurately predicted when user-agents behave differently. This may produce completed download results that are greater than 100%.
- Assets that take advantage of HTTP Progressive Download may not be accurately represented by this report. This is due to users seeking to different positions in a video.

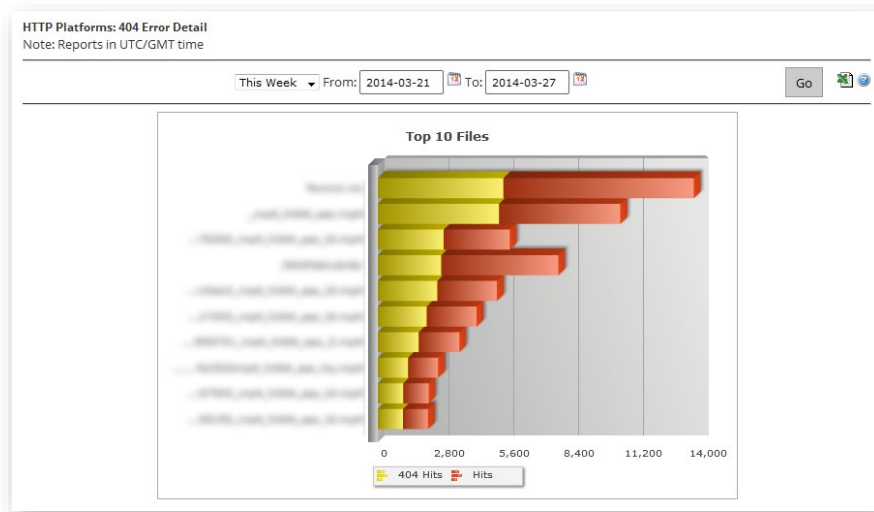
By 404 Errors Report

The By 404 Errors report allows you to identify the type of content that generates the most number of 404 Not Found status codes. The top of the report contains a bar chart for the top 10 assets for which a 404 Not Found status code was returned. This bar chart compares the total number of requests with requests that resulted in a 404 Not Found status code for those assets. Each bar is color-coded. A yellow bar is used to indicate that the request resulted in a 404 Not Found status code. A red bar is used to indicate the total number of requests for the asset.

Note: For the purposes of this report, a hit represents any request for an asset regardless of status code.

Note: This report is available for the HTTP Large and ADN platforms.

Reminder: For the purposes of this report, edge CNAME URLs are converted to their equivalent CDN URLs. This allows an accurate tally for all statistics associated with an asset regardless of the CDN or edge CNAME URL used to request it.



By 404 Errors Report for the HTTP Large platform

The left-hand side of the graph (y-axis) indicates the file name for each of the top 10 requested assets that resulted in a 404 Not Found status code. Directly below the graph (x-axis), you will find labels that indicate the total number of requests and the number of requests that resulted in a 404 Not Found status code.

Directly below the bar chart, the following information will be listed for the top 250 requested assets: relative path (including file name), the number of requests that resulted in a 404 Not Found status code, the total number of times that the asset was requested, and the percentage of requests that resulted in a 404 status code. A description is provided for each of these metrics in the **Appendix A: Advanced Content Analytics Fields**.

Real-Time Statistics

Overview

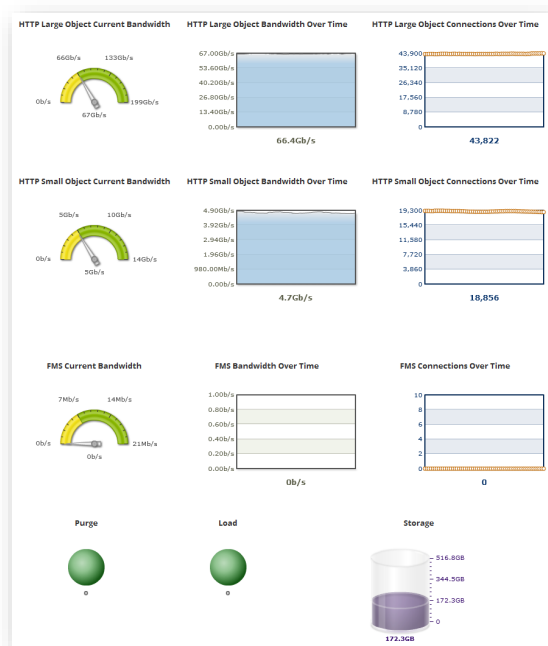
Real-Time Statistics provides real-time data about the performance of our CDN when delivering content to your clients. When viewing real-time statistics, you can choose to either view a statistical overview for all platforms or you can view more detailed information on a particular platform. Real-Time Statistics can be viewed by finding the **Analytics** menu and then selecting **Real-Time Stats**.

Note: Information on the Real-Time Alerts feature can be found in the **Real-Time Alerts** chapter.

Overview Report

The Overview report provides a single dashboard from which you can view real-time bandwidth and traffic data for all available platforms. Additionally, you can view the number of purge and load requests that are still in queue and the current amount of CDN storage space in use.

Note: If a particular platform has not been enabled on your account, then the corresponding graph(s) will not be available from the Overview report.




Overview Report for Real-Time Statistics

Current Bandwidth

A speedometer indicates current and historical bandwidth usage for each platform (e.g., HTTP Large or HTTP Small) that is active on your account. The current amount of bandwidth being used on a particular platform is displayed directly below the speedometer's needle. Additionally, each speedometer is color-coded to reflect historical data. The yellow portion of the speedometer indicates how much bandwidth was used over the last 24 hours. The green portion of the speedometer indicates how much bandwidth can be used without incurring overage charges. The exact upper-limit for bandwidth usage without incurring overage charges can be viewed on the right-hand side of the speedometer.

Note: The reported amount of bandwidth usage is rounded to the nearest whole number.

Note: The units used by to report bandwidth usage depends on the amount of traffic being served for your account on that platform and will be one of the following: bits per second (b/s), Kilobits per second (Kb/s), Megabits per second (Mb/s), or Gigabits per second (Gb/s).




Speedometer (Overview Report)

Bandwidth Over Time

The Bandwidth Over Time line graph provides a glimpse into recent bandwidth usage for each platform (e.g., HTTP Large or HTTP Small) that is active on your account. The exact amount of bandwidth being used on a particular platform is displayed directly below the line graph. A visual representation of this bandwidth usage is indicated by the shaded portion of the graph.

Note: The units used by to report bandwidth usage depends on the amount of traffic being served for your account on that platform and will be one of the following: bits per second (b/s), Kilobits per second (Kb/s), Megabits per second (Mb/s), or Gigabits per second (Gb/s).


New data on bandwidth usage is made available to this report every five seconds and then plotted on the Bandwidth Over Time line graph. This line graph contains approximately 5 minutes worth of data. This means that data that is older than 5 minutes will be removed from the graph to make room for newer data.



Bandwidth Over Time (Overview Report)

Connections Over Time

The Connections Over Time line graph reports the average number of new connections per second to our CDN service.



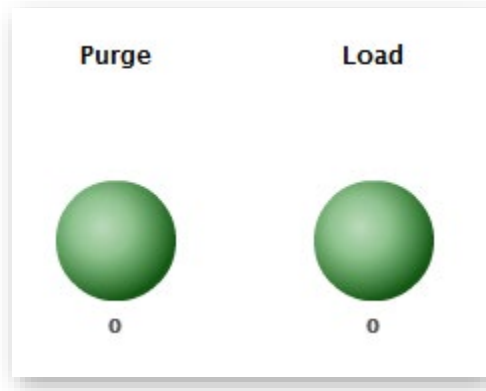
Connections Over Time (Overview Report)

Key information:

- This line graph contains approximately 5 minutes worth of data.
- New data on the number of new connections is plotted every five seconds. Data older than 5 minutes will be removed from the graph to make room for newer data.
- A user agent's (e.g., web browser) initial request for content will always establish a connection. After which, the user agent determines whether future requests within that session will reuse that connection or whether new connections will be established.
- This statistic is calculated using the following two steps:
 1. The average number of new connections per second on each edge server is calculated.
 2. This data is collected from all edge servers and then summed.

Purge and Load Requests

At the lower-left hand corner of the Overview report, you will find the number of purge and load requests that are in queue to be processed. You can view the exact number of purge/load requests under the respective sphere.




Purge & Load Queue (Overview Report)

Storage

The amount of CDN storage space being used can be viewed directly below the storage cylinder chart. The shaded portion of the cylinder indicates how much CDN storage space is in use. The portion that appears empty represents the amount of disk space that can be used without incurring overage charges. The exact upper-limit for storage disk space usage without incurring overage charges can be viewed on the upper right-hand side of the cylinder.

Note: CDN storage space usage is calculated once a day. This means that it may not reflect recent changes to your storage account.



CDN Storage Usage (Overview Report)

Detailed Real-Time Statistics

View detailed real-time statistics for a specific HTTP-based platform by clicking on it from the side navigation tab. The following types of real-time statistics will be shown:

- Bandwidth
- Status Codes
- Cache Statuses
- Connections

Life Span

The above graphs display real-time statistics for a set time period. A sliding window of data is displayed once the specified time has passed. This means that old data will be removed from the graph to make room for new data. Set the length of time for this sliding window via the **Time span of graphs** option.

Filters

Use the **Filtering Options** option to apply one of the following filters:


Filter	Description
CNAME	Filters all graphs for requests that point to the selected edge CNAME.
POP	Filters all graphs for requests that were processed by the selected POP.
Country	Filters all graphs for requests that were processed by a POP in the selected country.

Note: Only a single filter may be applied to these graphs at any given time.

Bandwidth

The Bandwidth graph displays the amount of bandwidth used for the current platform over a specified period of time. The shaded portion of the graph indicates bandwidth usage. The exact amount of bandwidth currently being used is displayed directly below the line graph.

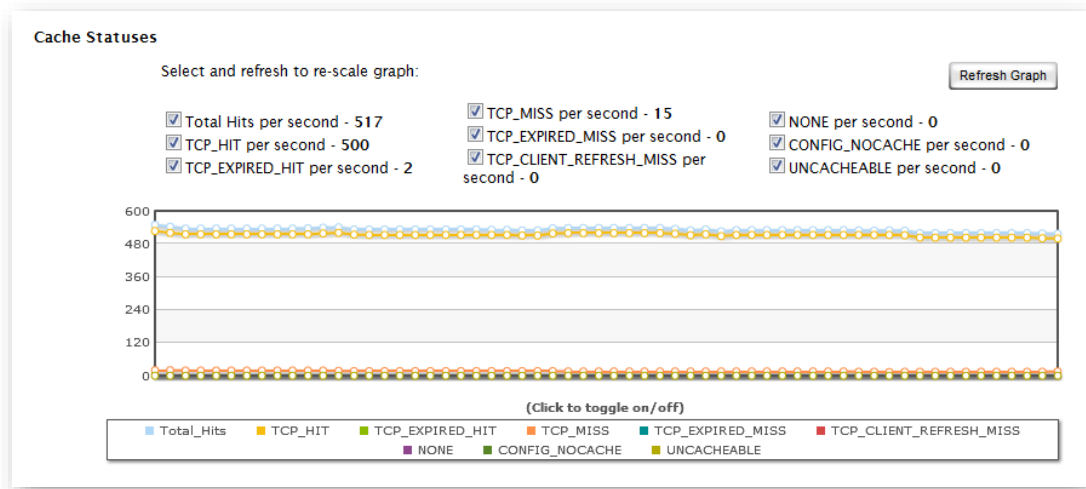
Note: The units used by to report bandwidth usage are one of the following: bits per second (b/s), Kilobits per second (Kb/s), Megabits per second (Mb/s), or Gigabits per second (Gb/s).



Real-Time Statistics (Bandwidth)

Status Codes

The Status Codes graph consists of color-coded lines that indicate how often HTTP response codes are occurring over a specified period of time. The left side of the graph (y-axis) indicates how often a status code is returned for requests, while the bottom of the graph (x-axis) indicates the progression of time.



Real-Time Statistics (Status Codes)

A list of status codes is displayed directly above the graph. This list indicates each status code that can be included in the line graph and the current number of occurrences per second for

that status code. By default, a line is displayed for each of these status codes in the graph. However, you can choose to only monitor the status codes that have special significance for your CDN configuration. This can be accomplished by only marking the desired status code options and clearing all other options. After you are satisfied with the status codes that will be displayed in the graph, you should click **Refresh Graph**. This will prevent the cleared status codes from being included in the graph.

Note: The **Refresh Graph** option will clear the graph. After which, it will only display the selected status codes.

Select and refresh to re-scale graph: Refresh Graph

<input checked="" type="checkbox"/> Total Hits per second - 520	<input checked="" type="checkbox"/> 3xx per second - 0	<input checked="" type="checkbox"/> 4xx per second - 0
<input checked="" type="checkbox"/> 2xx per second - 329	<input checked="" type="checkbox"/> 403 per second - 183	<input checked="" type="checkbox"/> 5xx per second - 0
<input checked="" type="checkbox"/> 304 per second - 7	<input checked="" type="checkbox"/> 404 per second - 1	<input checked="" type="checkbox"/> Other per second - 0

Status Code Options


Each status code option is described below.

Name	Description
Total Hits per second	Determines whether the total number of requests per second for the current platform will be displayed in the graph. You can use this option as a baseline indicator to see the percentage of total hits that a particular status code comprises.
2xx per second	Determines whether the total number of 2xx status codes (e.g., 200, 201, 202, etc.) that occur per second for the current platform will be displayed in the graph. This type of status code indicates that the request was successfully delivered to the client.
304 per second	Determines whether the total number of 304 status codes that occur per second for the current platform will be displayed in the graph. This status code indicates that the requested asset has not been modified, since it was last retrieved by the HTTP client.
3xx per second	Determines whether the total number of 3xx status codes (e.g., 300, 301, 302, etc.) that occur per second for the current platform will be displayed in the graph. This type of status code indicates that the request resulted in a redirection.
403 per second	Determines whether the total number of 403 status codes that occur per second for the current platform will be displayed in the graph. This status code indicates that the request was deemed unauthorized. One possible cause for this status code is when an unauthorized user requests an asset protected by Token-Based Authentication.

Name	Description
404 per second	Determines whether the total number of 404 status codes that occur per second for the current platform will be displayed in the graph. This status code indicates that the requested asset could not be found.
4xx per second	Determines whether the total number of 4xx status codes (e.g., 400, 401, 402, 405, etc.) that occur per second for the current platform will be displayed in the graph. This status code indicates that the requested asset was not delivered to the client.
5xx per second	Determines whether the total number of 5xx status codes (e.g., 500, 501, 502, etc.) that occur per second for the current platform will be displayed in the graph.
Other per second	Determines whether the total occurrences for all other status codes will be reported in the graph.

You can also choose to temporarily hide logged data for a particular status code. You may do this from the area directly below the graph by clearing the desired status code option. The selected status code will be immediately hidden from the graph. Marking that status code option will cause that option to be displayed again.

Note: The color-coded options directly below the graph only affect what is displayed in the graph. It does not affect whether the graph will keep track of that status code.




Showing/Hiding Status codes

Cache Statuses

The Cache Statuses graph consists of color-coded lines that indicate how often certain types of cache statuses are occurring over a specified period of time. The left side of the graph (y-axis) indicates how often a cache status is returned for requests, while the bottom of the graph (x-axis) indicates the progression of time.


Note: A definition is provided for each cache status from **Appendix C: Cache Statuses**.



Real-Time Statistics (Cache Status)

A list of cache statuses is displayed directly above the graph. This list indicates each cache status that can be included in the line graph and the current number of occurrences per second for that cache status. By default, a line is displayed for each of these cache statuses in the graph. However, you can choose to only monitor the cache statuses that have special significance for your CDN configuration. This can be accomplished by only marking the desired cache status options and clearing all other options. After you are satisfied with the cache statuses that will be displayed in the graph, you should click **Refresh Graph**. This will prevent the cleared status codes from being included in the graph.


Note: The **Refresh Graph** option will clear the graph. After which, it will only display the selected cache statuses.



Cache Status Options

You can also choose to temporarily hide logged data for a particular response code. You may do this by clearing the desired response code option from the area directly below the graph. The selected response code will be immediately hidden from the graph. Marking that response code option will cause that option to be displayed again.

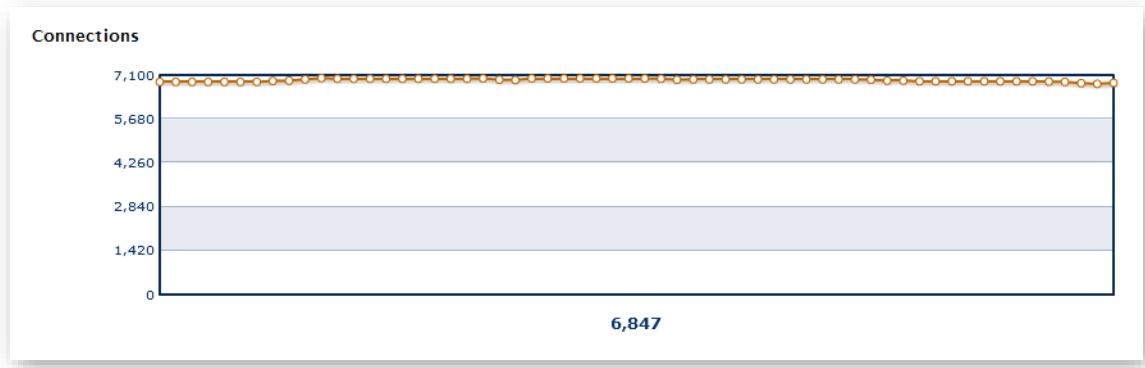
Note: The color-coded options directly below the graph only affect what is displayed in the graph. It does not affect whether the graph will keep track of that status code.



Showing/Hiding Response Codes

Connections

This line graph reports the average number of new connections per second to a specific delivery platform (i.e., HTTP Large, HTTP Small, or ADN).



Real-Time Statistics (Connections)

Key information:

- By default, this line graph contains approximately 5 minutes worth of data. Customize this window through the Time span of graphs option at the top of the page.
- New data on the number of new connections is plotted at regular intervals (e.g., five seconds when the graph's time span is 5 minutes). Data older than the time span defined in the **Time span of graphs** option will be removed from the graph to make room for newer data.
- A user agent's (e.g., web browser) initial request for content will always establish a connection. After which, the user agent determines whether future requests within that session will reuse that connection or whether new connections will be established.
- This statistic is calculated using the following two steps:
 1. The average number of new connections per second on each edge server is calculated.
 2. This data is collected from all edge servers and then summed.

Other Stats

The Other Stats report provides information over how many purge/load requests are currently being processed and the amount of CDN storage space being used.

Purge/Load Queue

Each request to purge or load an asset is placed in a queue. The number of assets that still need to be purged or loaded to our CDN is displayed directly below the corresponding sphere. If there are no items waiting to be processed, then "0" will be reported as the queue length.




Purge and Load Queue

Storage Usage

The Storage Usage portion of this report displays how much storage space is being used on our CDN origin servers. The exact amount of used disk space is displayed directly below the cylinder.

The portion of the cylinder that is filled in indicates how much storage space is in use. The portion that appears empty represents the amount of disk space that can be used without incurring overage charges. The exact upper-limit for storage disk space usage without incurring overage charges can be viewed on the upper right-hand side of the cylinder.



CDN Storage Space Utilization

Edge Performance Analytics

Overview

Edge Performance Analytics provides granular information on the amount of traffic and bandwidth for the HTTP Large, HTTP Small, and ADN platforms. This information can then be used to generate trending statistics, which allow you to gain insight on how your assets are being cached and delivered to your clients. In turn, this allows you to form a strategy on how to optimize the delivery of your content and to determine what issues should be tackled to better leverage the CDN. As a result, not only will you be able to improve data delivery performance, but you will also be able to reduce your CDN cost structure.

Note: All reports use UTC/GMT notation when specifying a date/time.

Reports & Log Collection

CDN activity data must be collected by the Edge Performance Analytics module before it can generate reports on it. This collection process occurs once a day and it covers the activity that took place during the previous day. This means that a report's statistics represent a sample of the day's statistics at the time it was processed, and do not necessarily contain the complete set of data for the current day. The primary function of these reports is to assess performance. They should not be used for billing purposes or exact numeric statistics.

Tip: If billing is your primary purpose for viewing a report, then you should generate and analyze reports (e.g., Traffic Summary, Bandwidth, or Data Transferred) in the Core Reporting module.

Note: The raw data from which Edge Performance Analytic reports are generated is available for at least 90 days.

Asset/Directory Location

Certain reports (e.g., By File, By File Detail, By Directory, etc.) need to specify a unique path to an asset or a directory. They are able to do so through the use of a CDN URL path. A CDN URL path starts with the content access point. This involves truncating the protocol and hostname from the CDN URL.

Note: For the purposes of the Edge Performance Analytics module, edge CNAME URLs are converted to their equivalent CDN URLs. This allows an accurate tally for all statistics associated with an asset regardless of the CDN or edge CNAME URL used to request it.

Sample CDN URL:

http://wac.0001.edgecastcdn.net/000001/main/index.html

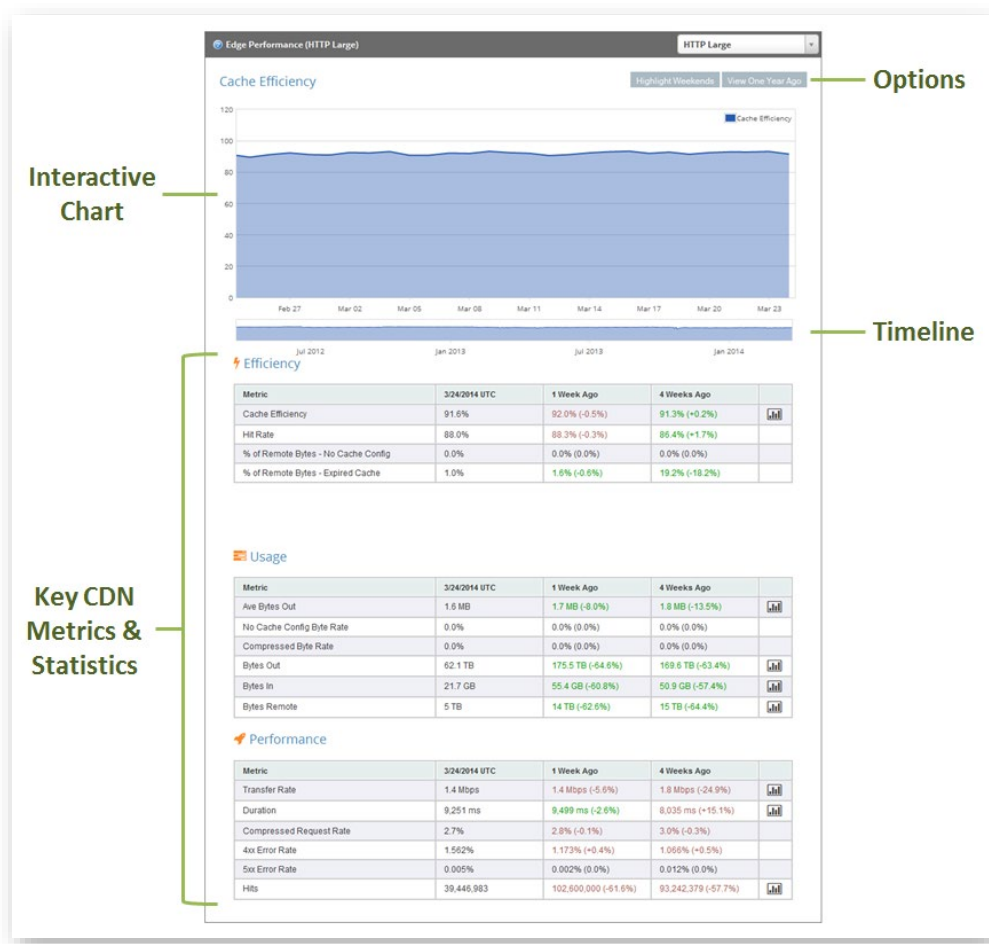
Sample CDN URL path:

/000001/main/index.html

The above sample CDN URL path starts with a content access point (e.g., /000001). A content access point (i.e., *yyxxxx* or *80xxxx/CustomerOrigin*) starts with an origin identifier. The first two digits of an origin identifier indicate whether the asset is stored on a CDN (i.e., 00) or a customer origin server (i.e., 80).

Dashboard

The Edge Performance Analytics dashboard tracks current and historical CDN traffic through a chart and statistics. Use this dashboard to detect recent and long-term trends on the performance of CDN traffic for your account.



Edge Performance Analytics Dashboard

This dashboard consists of:

- An interactive chart that allows the visualization of key metrics and trends.
- A timeline that provides a sense of long term patterns for key metrics and trends.
- Key metrics and statistical information on how our CDN network improves site traffic as measured by overall performance, cost, efficiency, and security.

Note: The upper-right hand corner of the dashboard contains a list of platforms. The dashboard will display data for the currently selected platform.

Note: The purpose of this dashboard is to identify areas where data delivery performance may be improved. This type of analysis can typically be performed by studying overall traffic patterns as opposed to measuring every request submitted to the CDN. Metrics that track traffic volume indicate the amount of the traffic used in the calculation of ratios and percentages, and may only show a portion of the total traffic for high-volume customers (approximately 50 million hits or more per day). Please leverage our Core Reporting module for precise traffic volume data.

Chart

The dashboard contains a chart that tracks a metric over the time period selected in the timeline that appears directly below it.




Edge Performance Analytics Dashboard Chart (Shown Here with the Hits Metric)

Key information:

- This chart is platform-specific.
- By default, the cache efficiency rate for the last 30 days will be charted.
- This chart is generated from data collated on a daily basis.
- Hovering over a day on the line graph will indicate a date and the value of the metric on that date.

- Click **Highlight Weekends** to toggle an overlay of light gray vertical bars that represent weekends onto the chart. This type of overlay is useful for identifying traffic patterns over weekends.
- Click **View One Year Ago** to toggle an overlay of the previous year's activity over the same time period onto the chart. This type of comparison provides insight into long-term CDN usage patterns. The upper-right hand corner of the chart contains a legend that indicates the color code for each line graph.


To update the chart

1. **Platform:** Chart data for a different platform by selecting it from the upper-right hand corner of the dashboard. Keep in mind that the entire dashboard will be updated to reflected data for the selected platform.
2. **Time Range:** Perform one of the following:
 - Select the desired region in the timeline. The chart will be updated with data that corresponds to the selected time period.
 - Double-click the chart to display all available historical data up to a maximum of two years.
3. **Metric:** Click the  (Generate Chart) icon that appears next to the desired metric. The chart and the timeline will be refreshed with data for the corresponding metric.

Note: The  (Generate Chart) icon indicates whether a metric may be charted.

Timeline

A timeline that graphs up to the last two years of CDN activity is displayed directly below the chart.

Note: Use the  (Generate Chart) icon to change the type of metric graphed in both the chart and the timeline.

This chart serves the following purposes:

- Reveals long term patterns and trends.
- Indicates the time frame being displayed in the chart. The highlighted section indicates the time frame charted above.
- Defines the time frame that will be displayed in the chart. Use one of the following methods:
 - Drag the cursor and select the desired time period.
 - Double-click the timeline to display all data in the above chart.

Key Metrics & Statistics

Key metrics and statistics on CDN activity are displayed directly below the chart/timeline. These key metrics are grouped according to how they affect the CDN experience. These categories are described below.

Category	Description
Efficiency	Provides a sense as to how much of your traffic is being cached on our network.
Performance	Analyze CDN performance metrics to detect suboptimal CDN configurations.
Usage	Analyze these metrics to assess how to reduce CDN expenditure.
Secure Traffic	Provides statistics on HTTPS traffic.

Deciphering Metrics & Statistics

The following information and statistics are reported for each metric:

Name	Report Date	Description
<i>Yesterday</i>	1 day ago	Reports the value of the metric for the previous day (UTC).
1 Week Ago	8 days ago	Reports the following information for the day that took place 1 week prior to yesterday (UTC): <ul style="list-style-type: none">• The value of the metric for that day.• The percentage of difference between yesterday (UTC) and 8 days ago (UTC).• The color codes used by this field are described below.<ul style="list-style-type: none">▪ Green: Indicates that yesterday's data was an improvement over the data from 8 days ago.▪ Red: Indicates that yesterday's data was a diminishment over the data from 8 days ago.▪ Black: Indicates that the metric in question is not applicable for CDN traffic on those dates.
4 Weeks Ago	29 days ago	Reports the following information for the day that took place 4 weeks prior to yesterday (UTC): <ul style="list-style-type: none">• The value of the metric for that day.• The percentage of difference between yesterday (UTC) and 29 days ago (UTC).• The color codes used by this field are described below.<ul style="list-style-type: none">▪ Green: Indicates that yesterday's data was an improvement over the data from 29 days ago.

Name	Report Date	Description
		<ul style="list-style-type: none"> ▪ Red: Indicates that yesterday's data was a diminishment over the data from 29 days ago. ▪ Black: Indicates that the metric in question is not applicable for CDN traffic on those dates.

Efficiency Metrics

The purpose of these metrics is to see whether cache efficiency can be improved. The main benefits derived from cache efficiency are:

- Reduced load on the origin server which may lead to:
 - Better web server performance.
 - Reduced operational costs.
- Improved data delivery acceleration since more requests will be served directly from the CDN.

Field	Description
Cache Efficiency	<p>Indicates the percentage of data transferred that was served from cache. This metric measures when a cached version of the requested content was served directly from the CDN (edge servers) to requesters (e.g., web browser).</p> <p>Expectation</p> <p>This metric should be 90% or higher.</p> <p>Interpretation & Remediation</p> <p>Improve cache efficiency by checking the following factors:</p> <ul style="list-style-type: none"> • Cache Policy: Frequent revalidations due to cache expiration may impact both data delivery performance and cache efficiency. <ul style="list-style-type: none"> ▪ Verify the cache policy defined for content that expires frequently. • HTTP Rules Engine: An origin server's cache policy can be overwritten through HTTP Rules Engine. <ul style="list-style-type: none"> ▪ Check for a rule that defines a restrictive cache policy. • Purging: Unnecessary purges force new requests to be forwarded to the origin server. This will generate unnecessary traffic from the origin server to the CDN. <ul style="list-style-type: none"> ▪ Verify that purges are targeted to the narrowest set of assets required to achieve your content freshness requirements. • Query Strings: Query-string caching diminishes cache efficiency by increasing the amount of cached content and cache misses. This is especially true when there are frequent requests with unique query

Field	Description
	<p>strings.</p> <ul style="list-style-type: none"> ▪ Review and adjust query string caching settings. The "standard-cache" configuration is optimal for cache efficiency. ▪ If query string caching is required, consider using HTTP Rules Engine to apply query string caching to the narrowest set of assets necessary. This can be accomplished by defining the desired type of request through match options and then defining the query string caching policy through the Cache-Key Query String feature. <ul style="list-style-type: none"> • Non-200 Responses: A high rate of non-200 responses require header data to be transferred from the origin server to the CDN. Statistical information on the HTTP status codes for CDN requests can be viewed in the Hits (Status Codes) report. <ul style="list-style-type: none"> ▪ Verify that all linked and referenced content can be found on the origin server. <p>This metric is calculated through the following formula:</p> <div style="background-color: #f0f0f0; padding: 5px; border: 1px solid #ccc;"> $\text{Cache Efficiency} = 100\% - ((\text{Bytes Remote}/\text{Bytes Out}) * 100)$ </div> <p>Historical data is color-coded as indicated below.</p> <ul style="list-style-type: none"> • Green: Indicates that the cache efficiency ratio increased yesterday when compared to historical data. • Red: Indicates that the cache efficiency ratio decreased yesterday when compared to historical data.
Hit Rate	<p>Indicates the percentage of requests that were served from cache. This metric measures when a cached version of the requested content was served directly from the CDN (edge servers) to requesters (e.g., web browser).</p> <p>Key information:</p> <p>This metric measures both of the following types of requests:</p> <ul style="list-style-type: none"> • Requests for fresh content. • Requests for stale content that were successfully revalidated. <p>Interpretation & Remediation</p> <p>Improve hit rate performance by following the remediation steps defined for the Cache Efficiency metric.</p> <p>This metric is calculated through the following formula:</p> <div style="background-color: #f0f0f0; padding: 5px; border: 1px solid #ccc;"> $((\# \text{ of Cache Hits} + \# \text{ of Expired Cache Hits}) / \text{Hits}) * 100$ </div> <p>Historical data is color-coded as indicated below.</p> <ul style="list-style-type: none"> • Green: Indicates that the percentage of cache hits increased yesterday

Field	Description
	<p>when compared to historical data.</p> <ul style="list-style-type: none"> • Red: Indicates that the percentage of cache hits decreased yesterday when compared to historical data.
<p>% of Remote Bytes – No Cache Config</p>	<p>Indicates the percentage of traffic that was served from origin servers to the CDN (edge servers) due to the Bypass Cache feature (HTTP Rules Engine).</p> <p>Interpretation & Remediation</p> <p>This metric only tracks origin server traffic that was not cached as a result of a rule that leverages the Bypass Cache feature. This feature will prevent caching regardless of the cache policy defined on the origin server.</p> <p>Improve this metric by checking the following factors:</p> <ul style="list-style-type: none"> • Scope: Review and adjust the rule's match criteria to the narrowest scope that will achieve the desired caching policy. • ADN: Consider migrating content that should not be cached to the ADN platform. The ADN platform specializes in the acceleration of this type of content. <p>This metric is calculated through the following formula:</p> <div style="background-color: #f0f0f0; padding: 5px; border: 1px solid #ccc;"> $\% \text{ of Remote Bytes - No Cache Config} = (\text{Remote Server Bypass Cache Traffic} / \text{Bytes Remote}) * 100$ </div> <p>Historical data is color-coded as indicated below.</p> <ul style="list-style-type: none"> • Green: Indicates that the rate of Bypass Cache traffic decreased yesterday when compared to historical data. • Red: Indicates that the rate of Bypass Cache traffic increased yesterday when compared to historical data.
<p>% of Remote Bytes – Expired Cache</p>	<p>Indicates the percentage of traffic that was served from origin servers to the CDN (edge servers) as a result of stale content revalidation.</p> <p>Interpretation & Remediation</p> <p>Improve metric performance by following the remediation steps defined for the Cache Efficiency metric.</p> <p>This metric is calculated through the following formula:</p> <div style="background-color: #f0f0f0; padding: 5px; border: 1px solid #ccc;"> $\% \text{ of Remote Bytes - Expired Cache} = (\text{Expired Cache Traffic} / \text{Bytes Out}) * 100$ </div> <p>Historical data is color-coded as indicated below.</p> <ul style="list-style-type: none"> • Green: Indicates that the rate of expired cache traffic decreased yesterday when compared to historical data. • Red: Indicates that the rate of expired cache traffic increased yesterday when compared to historical data.

Usage Metrics

The purpose of these metrics is to provide insight into the following cost-cutting measures:

- Minimizing operational costs through the CDN.
- Reducing CDN expenditures through cache efficiency and compression.

Field	Description
Ave Bytes Out	<p>Indicates the average number of bytes transferred for each request served from the CDN (edge servers) to the requester (e.g., web browser).</p> <p>Interpretation & Remediation</p> <p>Analyze this metric to determine whether traffic is being delivered over an optimal platform.</p> <ul style="list-style-type: none"> • High: A high value for this metric means that the HTTP Large platform can optimally deliver this type of traffic. • Low: A low value for this metric means that the HTTP Small platform can optimally deliver this type of traffic. <p>This metric is calculated through the following formula:</p> $\text{Ave Bytes Out} = \text{Bytes Out}/\text{Hits}$ <p>Historical data is color-coded as indicated below.</p> <ul style="list-style-type: none"> • Green: Indicates that the number of average bytes per request increased yesterday when compared to historical data. • Red: Indicates that the number of average bytes per request decreased yesterday when compared to historical data.
No Cache Config Byte Rate	<p>Indicates the percentage of traffic served from the CDN (edge servers) to the requester (e.g., web browser) that was not cached as a result of the Bypass Cache feature.</p> <p>Interpretation & Remediation</p> <p>Improve performance and efficiency by reducing this metric.</p> <ul style="list-style-type: none"> • Cache Policy: Check whether this type of content would benefit from caching. If so, adjust the rule to prevent the Bypass Cache feature from being applied to those assets. • ADN: Consider migrating this type of content to the ADN platform. The ADN platform specializes in the acceleration of this type of content. <p>This metric is calculated through the following formula:</p> $\text{No Cache Config Byte Rate} = 100\% - ((\text{Bypass Cache Response Bytes}/\text{Bytes Out}) * 100)$ <p>Historical data is color-coded as indicated below.</p>

Field	Description
	<ul style="list-style-type: none"> • Green: Indicates that the percentage of Bypass Cache traffic decreased yesterday when compared to historical data. • Red: Indicates that the percentage of Bypass Cache traffic increased yesterday when compared to historical data.
Compressed Byte Rate	<p>Indicates the percentage of traffic sent from the CDN (edge servers) to requesters (e.g., web browser) in a compressed format.</p> <p>Interpretation & Remediation</p> <p>This metric is affected by the following factors:</p> <ul style="list-style-type: none"> • Device: A shift in access device usage (e.g., desktop to mobile devices) may bypass compression policies. <ul style="list-style-type: none"> ▪ Consider optimizing the site's compression policy for mobile devices. • Missing Accept-Encoding: Requests for content that should be compressed are missing the Accept-Encoding header. This request header is required for compression. • Unsupported Compression Method: Requests that contain an Accept-Encoding request header that references an unsupported compression method will not be compressed. <ul style="list-style-type: none"> ▪ Origin Server Compression: Make sure that the Accept-Encoding request header uses a compression method supported by the origin server. ▪ Edge Server Compression: Supported compression methods are gzip, deflate, and bzip2. Make sure that the Accept-Encoding request header is set to one of those methods. • Non-compressible Content: Frequent requests for content that cannot be compressed. <p>This metric is calculated through the following formula:</p> <div style="background-color: #f0f0f0; padding: 5px; margin: 10px 0;"> $\text{Compressed Byte Rate} = 100\% - \left(\frac{\text{Uncompressed Bytes Out} + \text{Incompressible Bytes Out}}{\text{Bytes Out}} \right) * 100$ </div> <p>Historical data is color-coded as indicated below.</p> <ul style="list-style-type: none"> • Green: Indicates that the percentage of compressed bytes sent increased yesterday when compared to historical data. • Red: Indicates that the percentage of compressed bytes sent decreased yesterday when compared to historical data.

Field	Description
Bytes Out	<p data-bbox="467 247 1409 315">Indicates the amount of data, in bytes, that were delivered from the CDN (edge servers) to the requester (e.g., web browser).</p> <hr/> <p data-bbox="467 342 1409 443">Note: Traffic volume numbers represent traffic that was used in calculations of ratios and percentages, and may only show a portion of the total traffic for high-volume customers.</p> <hr/> <p data-bbox="467 485 672 514">Key information:</p> <ul data-bbox="516 537 1409 695" style="list-style-type: none"> <li data-bbox="516 537 1409 604">• This metric tracks data delivery for response headers and the response body. <li data-bbox="516 627 1409 695">• This metric may be used to calculate potential operational costs that were mitigated through CDN usage. <p data-bbox="467 716 829 745">Interpretation & Remediation</p> <p data-bbox="467 768 1024 798">This metric is affected by the following factors:</p> <ul data-bbox="516 821 1409 1062" style="list-style-type: none"> <li data-bbox="516 821 1409 888">• Eyeballs: This metric reflects changes in the number of users requesting content via the CDN. <li data-bbox="516 911 1409 1062">• File Size: This metric reflects changes in the average file size transferred through the CDN. <ul data-bbox="613 999 1409 1062" style="list-style-type: none"> <li data-bbox="613 999 1409 1062">▪ Consider reducing average file size by compressing content on either the origin server or edge servers. <p data-bbox="467 1085 1045 1115">Historical data is color-coded as indicated below.</p> <ul data-bbox="516 1138 1409 1295" style="list-style-type: none"> <li data-bbox="516 1138 1409 1205">• Green: Indicates that the number of bytes decreased yesterday when compared to historical data. <li data-bbox="516 1228 1409 1295">• Red: Indicates that the number of bytes increased yesterday when compared to historical data.
Bytes In	<p data-bbox="467 1314 1409 1381">Indicates the amount of data, in bytes, sent from requesters (e.g., web browser) to the CDN (edge servers).</p> <hr/> <p data-bbox="467 1409 1409 1509">Note: Traffic volume numbers represent traffic that was used in calculations of ratios and percentages, and may only show a portion of the total traffic for high-volume customers.</p> <hr/> <p data-bbox="467 1551 672 1581">Key information:</p> <p data-bbox="467 1604 1349 1671">This metric may be used to calculate potential operational costs that were mitigated through CDN usage.</p> <p data-bbox="467 1692 829 1722">Interpretation & Remediation</p> <p data-bbox="467 1745 1317 1812">The manner in which this metric is calculated varies by HTTP method as indicated below.</p> <ul data-bbox="516 1835 1036 1864" style="list-style-type: none"> <li data-bbox="516 1835 1036 1864">• GET/HEAD: It includes request headers.

Field	Description
	<ul style="list-style-type: none"> • POST: It includes both request headers and the request body. <p>Historical data is color-coded as indicated below.</p> <ul style="list-style-type: none"> • Green: Indicates that the number of bytes decreased yesterday when compared to historical data. • Red: Indicates that the number of bytes increased yesterday when compared to historical data.
Bytes Remote	<p>Indicates the amount of data, in bytes, sent from CDN and customer origin servers to the CDN (edge servers).</p> <hr/> <p>Note: Traffic volume numbers represent traffic that was used in calculations of ratios and percentages, and may only show a portion of the total traffic for high-volume customers.</p> <hr/> <p>Key information:</p> <p>This metric includes response headers and the response body sent from an origin server to edge servers.</p> <p>Interpretation & Remediation</p> <p>This metric is affected by the following factors:</p> <ul style="list-style-type: none"> • Updates: Frequent requests for content that has not been cached. <ul style="list-style-type: none"> ▪ Check whether popular content is constantly being updated. • Stale Content: Frequent requests for stale content triggering cache revalidation with the origin server. <ul style="list-style-type: none"> ▪ Check the cache policy for the content in question. Can the TTL be safely extended? • Cache Inefficiencies: Unusually high levels for this metric may be indicative of cache efficiency issues. <ul style="list-style-type: none"> ▪ Please check whether the Cache Efficiency metric is below normal levels. <p>Additional information on data sent from origin servers to edge servers.</p> <ul style="list-style-type: none"> • A GET request for content that has not been previously cached requires that the origin server return both response headers and a response body. • A GET request for stale content requires that the origin server return either of the following: <ul style="list-style-type: none"> ▪ New Version: If a new version of the requested content is found, then the origin server must return response headers and a response body. ▪ Same Version: If the same version of the requested content is

Field	Description
	<p>found, then the origin server will only return response headers.</p> <p>Historical data is color-coded as indicated below.</p> <ul style="list-style-type: none"> • Green: Indicates that the amount of remote bytes decreased yesterday when compared to historical data. • Red: Indicates that the amount of remote bytes increased yesterday when compared to historical data.

Performance Metrics

The purpose of these metrics is to track overall CDN performance for your traffic. Information on how to interpret these metrics for the purpose of improving CDN performance is provided below.

Field	Description
Transfer Rate	<p>Indicates the average rate at which content was transferred from the CDN to a requester.</p> <p>Interpretation & Remediation</p> <p>This metric is affected by the following factors:</p> <ul style="list-style-type: none"> • Proximity: Requesters that are located closer to a POP will typically experience faster transfer rates than those located further away. <ul style="list-style-type: none"> ▪ Low or reduced transfer rates may be indicative of a shift in where users are located. Is your company making gains in a new regional market? If so, consider switching traffic over to a premium network. Contact your CDN account manager to learn more. • Last Mile: The quality of the requester's last mile connection determines transfer rate performance. <ul style="list-style-type: none"> ▪ A shift in access device usage (e.g., desktop to mobile devices) can affect the quality of the last mile connection for those users. In turn, this may diminish transfer rate performance. Consider optimizing your site for mobile devices. • Internet Health: Internet health may impact transfer rate. For example, a fiber cut may diminish transfer rate performance. <p>This metric is calculated through the following formula:</p> <div style="background-color: #f0f0f0; padding: 5px; margin: 5px 0;"> $\text{Transfer Rate} = \text{Ave Bytes Out} / \text{Duration}$ </div> <p>Historical data is color-coded as indicated below.</p> <ul style="list-style-type: none"> • Green: Indicates that the average transfer rate increased yesterday when compared to historical data. • Red: Indicates that the average transfer rate decreased yesterday when

Field	Description
	compared to historical data.
Duration	<p>Indicates the average time, in milliseconds, it took to deliver an asset to a requester (e.g., web browser).</p> <p>Interpretation & Remediation</p> <p>This metric is affected by the following factors:</p> <ul style="list-style-type: none"> • File Size: This metric reflects changes in the average file size transferred through the CDN. <ul style="list-style-type: none"> ▪ Consider reducing average file size by compressing content on either the origin server or edge servers. • Proximity: Requesters that are located closer to a POP will typically experience shorter data transfer duration than those located further away. <ul style="list-style-type: none"> ▪ Higher average duration may be indicative of a shift in where users are located. Is your company making gains in a new regional market? If so, consider switching traffic over to a premium network. Contact your CDN account manager to learn more. • Last Mile: The quality of the requester's last mile connection affects duration. <ul style="list-style-type: none"> ▪ A shift in access device usage (e.g., desktop to mobile devices) can affect the quality of the last mile connection for those users. In turn, this may diminish the data transfer duration. Consider optimizing your site for mobile devices. • Internet Health: Internet health may impact duration. For example, a fiber cut may diminish data transfer duration. <p>Check for the following patterns by comparing yesterday's data to data collected from 1 and 4 weeks ago. Look for significant deviations.</p> <ul style="list-style-type: none"> • An increase of this metric while the Transfer Rate metric remains level is an indicator of an increase in the average file size. • An increase of this metric while the Compressed Byte Rate metric decreases is an indicator that compression settings require adjustment. • An increase in both this metric and the % of Remote Bytes - No Cache Config metric is an indicator that the no-cache setting requires adjustment. <p>Historical data is color-coded as indicated below.</p> <ul style="list-style-type: none"> • Green: Indicates that the average duration decreased yesterday when compared to historical data. • Red: Indicates that the average duration increased yesterday when

Field	Description
	compared to historical data.
Compressed Request Rate	<p data-bbox="467 304 1328 367">Indicates the percentage of hits that were delivered from the CDN (edge servers) to the requester (e.g., web browser) in a compressed format.</p> <p data-bbox="467 388 829 420">Interpretation & Remediation</p> <p data-bbox="467 441 1023 472">This metric is affected by the following factors:</p> <ul data-bbox="516 493 1404 1144" style="list-style-type: none"> <li data-bbox="516 493 1404 598">• Device: A shift in access device usage (e.g., desktop to mobile devices) may bypass compression policies. Consider optimizing your site for mobile devices. <li data-bbox="516 619 1404 682">• Missing Accept-Encoding: Requests are missing the Accept-Encoding header. This request header is required for compression. <li data-bbox="516 703 1404 1060">• Unsupported Compression Method: Requests that contain an Accept-Encoding request header that is set to an unsupported compression method will result in the delivery of uncompressed data. <ul data-bbox="609 829 1404 1060" style="list-style-type: none"> <li data-bbox="609 829 1404 934">▪ Origin Server Compression: Make sure that the Accept-Encoding request header uses a compression method supported by the origin server. <li data-bbox="609 955 1404 1060">▪ Edge Server Compression: Supported compression methods are gzip, deflate, and bzip2. Make sure that the Accept-Encoding request header is set to one of those methods. <li data-bbox="516 1081 1404 1144">• Non-compressible Content: Frequent requests for content that cannot be compressed. <p data-bbox="467 1165 1047 1197">Historical data is color-coded as indicated below.</p> <ul data-bbox="516 1218 1328 1375" style="list-style-type: none"> <li data-bbox="516 1218 1328 1281">• Green: Indicates that the rate of compressed requests increased yesterday when compared to historical data. <li data-bbox="516 1302 1328 1375">• Red: Indicates that the rate of compressed requests decreased yesterday when compared to historical data.

Field	Description
4xx Error Rate	<p data-bbox="467 247 1243 279">Indicates the percentage of hits that generated a 4xx status code.</p> <p data-bbox="467 300 829 331">Interpretation & Remediation</p> <p data-bbox="467 352 1263 384">The following factors may increase the frequency of 4xx responses:</p> <ul data-bbox="516 405 1406 1035" style="list-style-type: none"> <li data-bbox="516 405 1406 531">• Missing Content: Frequent requests for content that is missing from an origin server. <ul style="list-style-type: none"> <li data-bbox="613 489 1049 520">▪ Identify and update invalid URLs. <li data-bbox="516 541 1406 783">• Invalid Cached Content: Frequent requests for cached pages that contain errors, such as incorrect URLs. <ol style="list-style-type: none"> <li data-bbox="613 636 1300 741">1. Verify that content on the origin server contains valid hyperlinks (i.e., href attribute) and references (i.e., src attribute). <li data-bbox="613 751 1130 783">2. Purge cached pages that contain errors. <li data-bbox="516 804 1406 1035">• Inline Linking: Sites that leverage inline linking (a.k.a. hotlinking or leeching) can quickly become obsolete when the referenced content is removed from your origin server. <ul style="list-style-type: none"> <li data-bbox="613 930 1406 1035">▪ Check for and prevent inline linking. Unauthorized links to content may be prevented by securing it through Token-Based Authentication and/or HTTP Rules Engine. <p data-bbox="467 1056 1049 1087">Historical data is color-coded as indicated below.</p> <ul data-bbox="516 1108 1406 1266" style="list-style-type: none"> <li data-bbox="516 1108 1406 1182">• Green: Indicates that the rate of 4xx errors decreased yesterday when compared to historical data. <li data-bbox="516 1192 1406 1266">• Red: Indicates that the rate of 4xx errors increased yesterday when compared to historical data.

Field	Description
5xx Error Rate	<p data-bbox="467 254 1243 279">Indicates the percentage of hits that generated a 5xx status code.</p> <p data-bbox="467 302 829 327">Interpretation & Remediation</p> <p data-bbox="467 352 1263 378">The following factors may increase the frequency of 5xx responses:</p> <ul data-bbox="516 407 1409 768" style="list-style-type: none"> <li data-bbox="516 407 1409 646"> <p data-bbox="516 407 1409 432">Origin Server: A non-responsive origin server may cause 5xx responses.</p> <ul data-bbox="610 457 1377 646" style="list-style-type: none"> <li data-bbox="610 457 1377 520">Verify that the origin server is responding properly by submitting a request to it. <li data-bbox="610 541 1377 646">If the origin server is being load balanced, then the 5xx issue may only arise when the request is load balanced to the symptomatic server. <li data-bbox="516 667 1370 768"> <p data-bbox="516 667 1370 768">Denial of Service: This type of response may be a symptom that the origin server is undergoing a DoS attack. Identify and mitigate/remediate malicious attacks on an origin server.</p> <p data-bbox="467 793 1045 819">Historical data is color-coded as indicated below.</p> <ul data-bbox="516 844 1396 995" style="list-style-type: none"> <li data-bbox="516 844 1396 907"> <p data-bbox="516 844 1396 907">Green: Indicates that the rate of 5xx errors decreased yesterday when compared to historical data.</p> <li data-bbox="516 928 1360 995"> <p data-bbox="516 928 1360 995">Red: Indicates that the rate of 5xx errors increased yesterday when compared to historical data.</p>

Field	Description
Hits	<p data-bbox="467 247 1068 279">Indicates the number of requests for CDN content.</p> <hr data-bbox="467 296 1417 300"/> <p data-bbox="467 306 1406 407">Note: Traffic volume numbers represent traffic that was used in calculations of ratios and percentages, and may only show a portion of the total traffic for high-volume customers.</p> <hr data-bbox="467 413 1417 417"/> <p data-bbox="467 449 672 480">Key information:</p> <ul data-bbox="516 504 1417 898" style="list-style-type: none"> <li data-bbox="516 504 1162 535">• A hit represents a request for content via the CDN. <ul data-bbox="610 554 1417 779" style="list-style-type: none"> <li data-bbox="610 554 1417 655">▪ A single hit would be logged for the following sample request: http://wpc.0001.edgecastcdn.net/000001/marketing/brochure.pdf <li data-bbox="610 674 1417 779">▪ Typically, a request for an asset generates a single hit. However, dynamic streaming technologies generate a multitude of hits for a single live stream or on-demand content. <li data-bbox="516 800 1390 898">• All requests for CDN content are included in this metric. This means that a hit is logged regardless of the HTTP status code returned to the requester. <p data-bbox="467 924 829 955">Interpretation & Remediation</p> <p data-bbox="467 974 1024 1005">This metric is affected by the following factors:</p> <ul data-bbox="516 1029 1417 1518" style="list-style-type: none"> <li data-bbox="516 1029 1417 1094">• Eyeballs: This metric reflects changes in the number of users requesting content via the CDN. <li data-bbox="516 1115 1417 1304">• Content: This metric reflects changes in the number of assets being requested via the CDN. <ul data-bbox="610 1203 1417 1304" style="list-style-type: none"> <li data-bbox="610 1203 1417 1304">▪ Consider whether consolidating multiple assets into a single asset would improve site performance by reducing the number of requests that need to be fulfilled. <li data-bbox="516 1325 1417 1430">• Popularity: A popular web page with 40 supporting assets (e.g., css, js, images, etc.) may generate more hits than a less popular web page with 80 assets. <li data-bbox="516 1451 1382 1518">• Dynamic Streaming: A single viewer may generate hundreds or even thousands of hits for a single live stream or on-demand video. <p data-bbox="467 1539 1045 1570">Historical data is color-coded as indicated below.</p> <ul data-bbox="516 1591 1360 1745" style="list-style-type: none"> <li data-bbox="516 1591 1360 1656">• Green: Indicates that the number of hits increased yesterday when compared to historical data. <li data-bbox="516 1677 1341 1745">• Red: Indicates that the number of hits decreased yesterday when compared to historical data.

Secure Traffic Metrics

The purpose of these metrics is to track CDN performance for HTTPS traffic. Information on how to interpret these metrics in an effort to improve CDN performance is provided below.

Field	Description
Secure Cache Efficiency	<p data-bbox="467 407 1349 541">Indicates the percentage of data transferred for HTTPS requests that were served from cache. This metric measures when a cached version of the requested content was served directly from the CDN (edge servers) to requesters (e.g., web browser) over HTTPS.</p> <p data-bbox="467 564 610 594">Expectation</p> <p data-bbox="467 617 899 646">This metric should be 90% or higher.</p> <p data-bbox="467 669 829 699">Interpretation & Remediation</p> <p data-bbox="467 722 1398 785">Improve secure cache efficiency by applying the remediation steps defined for the Cache Efficiency metric to HTTPS requests.</p> <p data-bbox="467 808 1398 905">Compare this metric to the Cache Efficiency metric. A significant difference between these two metrics is indicative of a configuration difference between secure and non-secure traffic.</p> <ul data-bbox="516 928 1089 957" style="list-style-type: none">• Verify the caching policy for HTTPS requests. <p data-bbox="467 980 1122 1010">This metric is calculated through the following formula:</p> <div data-bbox="477 1033 1419 1119" style="background-color: #f0f0f0; padding: 5px;">$\text{Secure Cache Efficiency} = 100\% - ((\text{Remote HTTPS Bytes}/\text{HTTPS Bytes Out}) * 100)$</div> <p data-bbox="467 1125 1045 1155">Historical data is color-coded as indicated below.</p> <ul data-bbox="516 1178 1325 1331" style="list-style-type: none">• Green: Indicates that the secure cache efficiency ratio increased yesterday when compared to historical data.• Red: Indicates that the secure cache efficiency ratio decreased yesterday when compared to historical data.

Field	Description
Secure Transfer Rate	<p>Indicates the average rate at which content was transferred from the CDN (edge servers) to requesters (e.g., web servers) over HTTPS.</p> <p>Interpretation & Remediation</p> <p>Analyze and improve secure transfer rate by applying the remediation steps defined for the Transfer Rate metric to HTTPS requests.</p> <p>This metric is calculated through the following formula:</p> <div style="background-color: #f0f0f0; padding: 5px; margin: 10px 0;"> $\text{Secure Transfer Rate} = \text{Secure Bytes Out} / \text{Average Secure Duration}$ </div> <p>Historical data is color-coded as indicated below.</p> <ul style="list-style-type: none"> • Green: Indicates that the average transfer rate for HTTPS requests increased yesterday when compared to historical data. • Red: Indicates that the average transfer rate for HTTPS requests decreased yesterday when compared to historical data.
Average Secure Duration	<p>Indicates the average time, in milliseconds, it took to deliver an asset to a requester (e.g., web browser) over HTTPS.</p> <p>Interpretation & Remediation</p> <p>Analyze and improve secure duration by applying the remediation steps defined for the Duration metric to HTTPS requests.</p> <p>Check for the following patterns by comparing yesterday's data to data collected from 1 and 4 weeks ago. Look for significant deviations.</p> <ul style="list-style-type: none"> • Average Secure Transfer Rate: An increase in duration while the transfer rate remains level is an indicator of an increase in the average file size. • Secure Cache Efficiency: Elevated levels of duration may be due to inefficient caching. The average time it takes to deliver content is suboptimal when content is being served from origin servers instead of directly from the edge of our network. <p>Historical data is color-coded as indicated below.</p> <ul style="list-style-type: none"> • Green: Indicates that the average duration for HTTPS requests decreased yesterday when compared to historical data. • Red: Indicates that the average duration for HTTPS requests increased yesterday when compared to historical data.

Field	Description
Secure Hits	<p data-bbox="464 252 1149 279">Indicates the number of HTTPS requests for CDN content.</p> <hr/> <p data-bbox="464 306 1403 407">Note: Traffic volume numbers represent traffic that was used in calculations of ratios and percentages, and may only show a portion of the total traffic for high-volume customers.</p> <hr/> <p data-bbox="464 453 672 480">Key information:</p> <ul data-bbox="516 506 1419 831" style="list-style-type: none"> <li data-bbox="516 506 1419 743">• A hit represents an HTTPS request for content via the CDN. <ul data-bbox="610 558 1419 743" style="list-style-type: none"> <li data-bbox="610 558 1419 659">▪ A single hit would be logged for the following sample request: https://wpc.0001.edgecastcdn.net/000001/marketing/brochure.pdf <li data-bbox="610 680 1419 743">▪ Typically, a request for an asset generates a single hit. However, dynamic streaming technologies generate a multitude of hits. <li data-bbox="516 768 1419 831">• Requests for CDN content, regardless of the HTTP status code returned to the requester, are tracked by this metric. <p data-bbox="464 856 829 884">Interpretation & Remediation</p> <p data-bbox="464 909 1024 936">This metric is affected by the following factors:</p> <ul data-bbox="516 961 1419 1444" style="list-style-type: none"> <li data-bbox="516 961 1419 1024">• Eyeballs: This metric reflects changes in the number of users requesting content via the CDN. <li data-bbox="516 1045 1419 1234">• Content: This metric reflects changes in the number of assets being requested via the CDN. <ul data-bbox="610 1136 1419 1234" style="list-style-type: none"> <li data-bbox="610 1136 1419 1234">▪ Consider whether consolidating multiple assets into a single asset would improve site performance by reducing the number of requests that need to be fulfilled. <li data-bbox="516 1255 1419 1356">• Popularity: A popular web page with 40 supporting assets (e.g., css, js, images, etc.) may generate more hits than a less popular web page with 80 assets. <li data-bbox="516 1377 1419 1444">• Dynamic Streaming: A single viewer may generate hundreds or even thousands of hits for a single video. <p data-bbox="464 1470 1045 1497">Historical data is color-coded as indicated below.</p> <ul data-bbox="516 1522 1419 1675" style="list-style-type: none"> <li data-bbox="516 1522 1419 1585">• Green: Indicates that the number of secure requests increased yesterday when compared to historical data. <li data-bbox="516 1606 1419 1675">• Red: Indicates that the number of secure requests decreased yesterday when compared to historical data.

Field	Description
Secure Bytes Out	<p data-bbox="464 258 1377 317">Indicates the amount of HTTPS traffic, in bytes, that were delivered from the CDN (edge servers) to the requester (e.g., web browser).</p> <hr/> <p data-bbox="464 342 1403 443">Note: Traffic volume numbers represent traffic that was used in calculations of ratios and percentages, and may only show a portion of the total traffic for high-volume customers.</p> <hr/> <p data-bbox="464 489 672 514">Key information:</p> <ul data-bbox="516 539 1403 695" style="list-style-type: none"> <li data-bbox="516 539 1403 606">• This metric tracks data delivery for response headers and the response body. <li data-bbox="516 632 1403 695">• This metric may be used to calculate potential operational costs that were mitigated through CDN usage. <p data-bbox="464 720 829 745">Interpretation & Remediation</p> <p data-bbox="464 770 1019 795">This metric is affected by the following factors:</p> <ul data-bbox="516 821 1414 1060" style="list-style-type: none"> <li data-bbox="516 821 1414 888">• Eyeballs: This metric reflects changes in the number of users requesting content via the CDN. <li data-bbox="516 913 1414 1060">• File Size: This metric reflects changes in the average file size transferred through the CDN. <ul data-bbox="610 997 1398 1060" style="list-style-type: none"> <li data-bbox="610 997 1398 1060">▪ Consider reducing average file size by compressing content on either the origin server or edge servers. <p data-bbox="464 1085 1045 1110">Historical data is color-coded as indicated below.</p> <ul data-bbox="516 1136 1398 1291" style="list-style-type: none"> <li data-bbox="516 1136 1398 1203">• Green: Indicates that the amount of secure content requested yesterday decreased when compared to historical data. <li data-bbox="516 1228 1398 1291">• Red: Indicates that the amount of secure content requested yesterday increased when compared to historical data.

Scheduled E-mails and Alerts

Report data for Edge Performance Analytics can either be:

- Scheduled for automated delivery to one or more email addresses
- Sent as an alert to one or more email addresses when a certain criteria is met.

Scheduled Report Delivery

Automated e-mail delivery of Edge Performance Analytics reports may be scheduled according to the frequency defined below.

Frequency	Description	Delivery
Daily	Reports the data collected for the previous day (GMT).	Approximately 48 Hours
Weekly	Reports the data collected for the previous week from Saturday (GMT) to Friday (GMT).	Monday
Monthly	Reports the data for the previous month (GMT).	Second Day of the Month

Automated report delivery allows quick and consistent review of CDN activity without requiring manual report generation. In order to maximize readability across multiple email clients, the data for a particular report is displayed in table format.

Tip: Report data may also be included as a CSV file attachment by enabling the **Attach report as CSV file** option when setting up its delivery.

Note: Reports delivered by email exclude graphs. View the graph for a particular report by generating the same report within the MCC.

To schedule e-mail reports

1. Navigate to the **Edge Performance Analytics** page. Load this page by finding the **Analytics** menu and then selecting **Edge Performance Analytics**.
2. Expand the desired platform and then select the desired report.
3. Click the **Schedule E-mails / Alerts** link.
4. Make sure that the **Schedule E-mail** option is selected.
5. From the **Frequency** option, select whether a report will be sent on a daily, weekly, or monthly basis. This option also determines the time period that the report will cover.
6. In the **To E-mail Addresses (comma delimited)** option, type one or more e-mail addresses to which the report will be sent.

Tip: Use a comma to delimit multiple email addresses (e.g., administrators@example.com,IT@example.com).

7. In the **Start Date** option, specify the initial date (YYYY-MM-DD) on which e-mails will be sent out. If you would like to consult a calendar, click the calendar icon that appears to the right of this option.
8. If you only want to receive reports by e-mail until a certain date, then you should specify that date (YYYY-MM-DD) in the **End Date (optional)** option. Otherwise, if you would like to receive reports via e-mail indefinitely, then you should leave this option blank.
9. Determine whether report data will be attached to the email as a CSV file by toggling the **Attach report as CSV file** option.
10. Make sure that the **Enabled** option is selected. Clearing this option will disable this configuration from sending reports by e-mail.
11. Click **Add** to save your configuration.

Alerts

In addition to receiving reports through e-mail, you may also be alerted as to when a particular condition arises. This allows you to catch problems quicker and therefore solve them sooner. An alert can be configured for requests denied due to Token-Based Authentication or any of the following response codes: 403 Forbidden, 404 Not Found, 4xx, 502 Bad Gateway, 504 Gateway Timeout, and 5xx.

When specifying an alert, you will need to specify the threshold that must be passed before the alert is sent out. This process involves selecting the criteria for determining when an alert should be sent, a mathematical operator, and the value that must be met. The available operators are defined below.

Operator	Description
<	This operator indicates that the percentage or number of occurrences must be less than the specified amount.
<=	This operator indicates that the percentage or number of occurrences must be less than or equal to the specified amount.
>	This operator indicates that the percentage or number of occurrences must be greater than the specified amount.
>=	This operator indicates that the percentage or number of occurrences must be greater than or equal to the specified amount.
=	This operator indicates that the percentage or number of occurrences must be equal to the specified amount.

An example of an alert is to configure an e-mail notification for when the percentage of 404 Not Found response codes being returned in any given day is greater than 10% (i.e., % 404 Hits > 10).

Tip: In order to simplify e-mail management, we suggest that you use e-mail distribution lists for alerts and reports. This allows you to simply modify a distribution list when adding or removing e-mail addresses.

Tip: Alert data may also be included as a CSV file attachment by enabling the **Attach report as CSV file** option when setting up its delivery.

Note: Each alert is based on the information gathered for the current day. For the purposes of alerts and reports, a day starts at midnight (i.e., 00:00 GMT).

To configure an e-mail alert

1. Navigate to the **Edge Performance Analytics** page. Load this page by finding the **Analytics** menu and then selecting **Edge Performance Analytics**.
2. Expand the desired platform and then select one of the following reports: Token Auth Deny Details, 404 Errors, 403 Errors, 4xx Errors, 504 Errors, 502 Errors, or 5xx Errors.
3. Click the **Schedule E-mails / Alerts** link.
4. Select the **Schedule Alert** option.
5. In the **Alert Threshold** option, perform the following:
 - i. Select whether an alert will be sent out based on the number or percentage of occurrences.
 - ii. Select the mathematical operator that will be used to determine when an alert should be sent out.
 - iii. Specify the threshold value for sending an e-mail notification. This threshold value will either be a percentage or the number of occurrences of the selected error code. Make sure that you do not specify units when specifying this threshold value.
6. In the **To E-mail Addresses (comma delimited)** option, type one or more e-mail addresses to which the report will be sent. When specifying multiple e-mail addresses, use a comma to delimit each one (e.g., administrators@domain.com,IT@domain.com).
7. In the **Start Date** option, specify the initial date (YYYY-MM-DD) on which e-mails will be sent out. If you would like to consult a calendar, click the calendar icon that appears to the right of this option.
8. If you only want to receive reports by e-mail until a certain date, then you should specify that date (YYYY-MM-DD) in the **End Date (optional)** option. Otherwise, if you would like to receive reports by e-mail indefinitely, then you should leave this option blank.
9. Determine whether alert data will be attached to the email as a CSV file by toggling the **Attach report as CSV file** option.
10. Make sure that the **Enabled** option is selected. Clearing this option will disable this configuration from sending reports by e-mail.
11. Click **Add** to save your configuration.

Reports

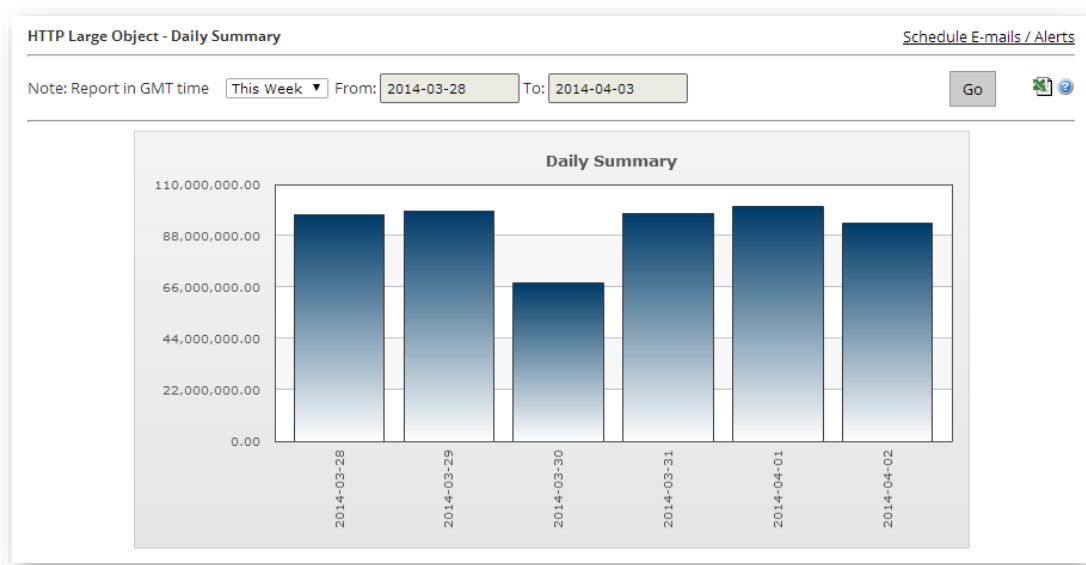
This section discusses the types of Edge Performance Analytics that can be generated for the HTTP Large, HTTP Small, and ADN platforms. For each of these reports, a chart allows you to quickly assess data trends. Directly below this chart, you will find a table containing pertinent bandwidth and traffic information. This information can be used to delve deeper into how your content is being served to your clients.

Note: The time chunk interval for all Edge Performance Analytics reports is daily. Report data will always include the specified end date.

Daily Summary Report

The Daily Summary report allows you to view daily traffic trends over a specified time period. Each bar on this graph represents a particular date. The size of the bar indicates the total number of hits that occurred on that date.

Tip: Hover the cursor over a bar to view the date (YYYY-MM-DD) corresponding to it and the total number of hits that occurred on that date.



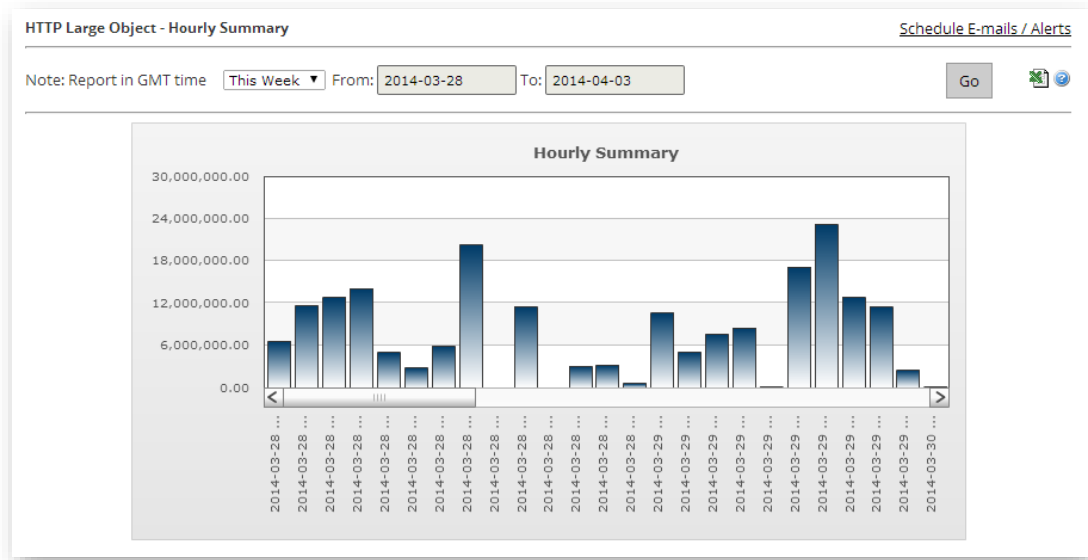
Daily Summary Report for the HTTP Large Platform

The data that was used to generate the chart, along with detailed bandwidth and traffic information, is displayed directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

Hourly Summary Report

The Hourly Summary report allows you to view hourly traffic trends over a specified time period. Each bar on this graph represents a single hour on a particular date. The size of the bar indicates the total number of hits that occurred during that hour.

Tip: Hover the cursor over a bar to view the date and time (YYYY-MM-DD hh:mm) corresponding to it and the total number of hits that occurred during the specified hour.



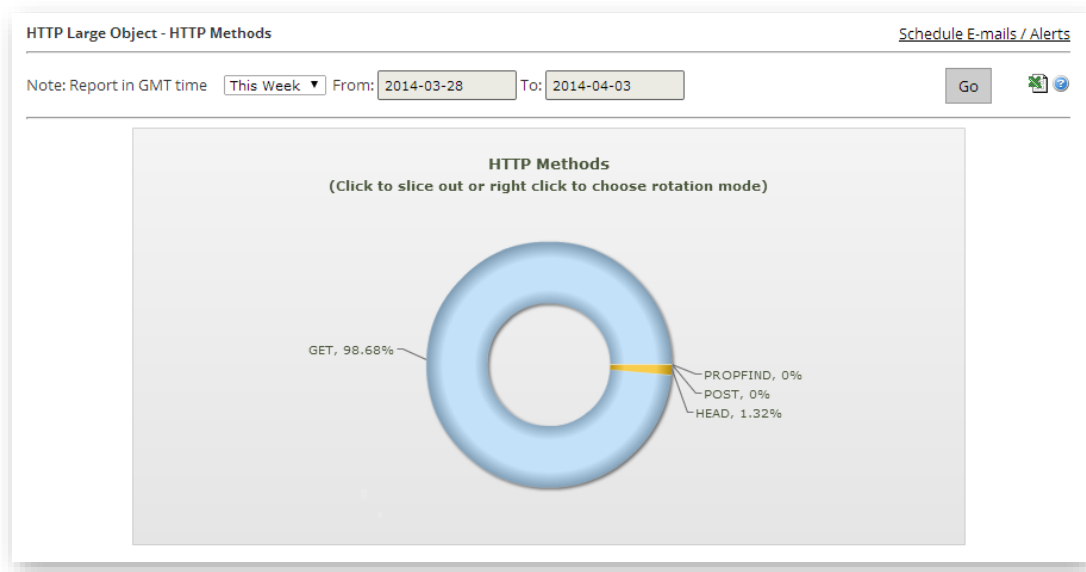
Hourly Summary Report for the HTTP Large Platform

The data that was used to generate the chart, along with detailed bandwidth and traffic information, is displayed directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

Protocols Report

The Protocols report displays the breakdown of traffic between the HTTP and HTTPS protocols. A donut chart indicates the percentage of hits that occurred for each type of protocol.

Note: If you have not purchased the SSL Traffic feature, then this report will indicate that all of your traffic used the HTTP protocol.

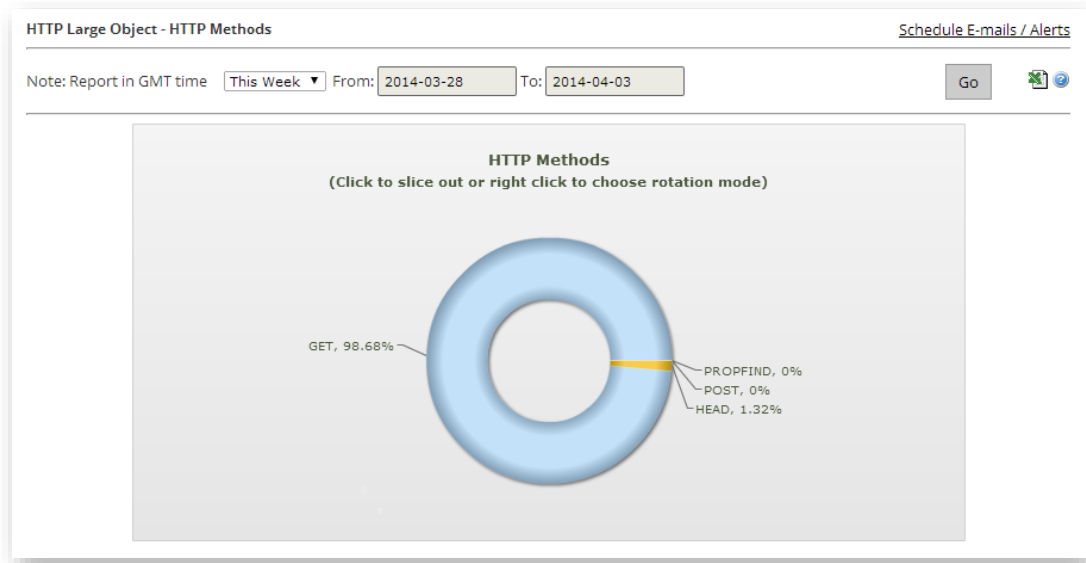


Protocols Report for the HTTP Large Platform

The data that was used to generate the chart, along with detailed bandwidth and traffic information, is displayed directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

HTTP Methods Report

The HTTP Methods report allows you to get a quick sense of which HTTP methods are being used to request your data. Typically, the most common HTTP request methods are GET, HEAD, and POST. A donut chart indicates the percentage of hits that occurred for each type of HTTP request method.




HTTP Methods Report for the HTTP Large Platform

The data that was used to generate the chart, along with detailed bandwidth and traffic information, is displayed directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

URLs Report

The URLs report contains a graph that displays the top 10 requested URLs. A bar is displayed for each URL. The height of the bar indicates how many hits that particular URL has generated over the time span covered by the report. Statistics for the top 100 requested URLs are displayed directly below this graph.

Tip: Hover the cursor over a bar to view the relative CDN URL path for the asset associated with it.




URLs Report for the HTTP Large Platform (Statistics Shown for the Top 10 Requested URLs)

The data that was used to generate the chart, along with detailed bandwidth and traffic information, is displayed directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

Cnames Report

The CNAMEs report contains a graph that displays the top 10 CNAMEs used to request assets over the time span of a report. Statistics for the top 100 requested CNAMEs are displayed directly below this graph.

Note: Variations in the domain (i.e., due to case or port) will cause a separate entry to be listed in this field.



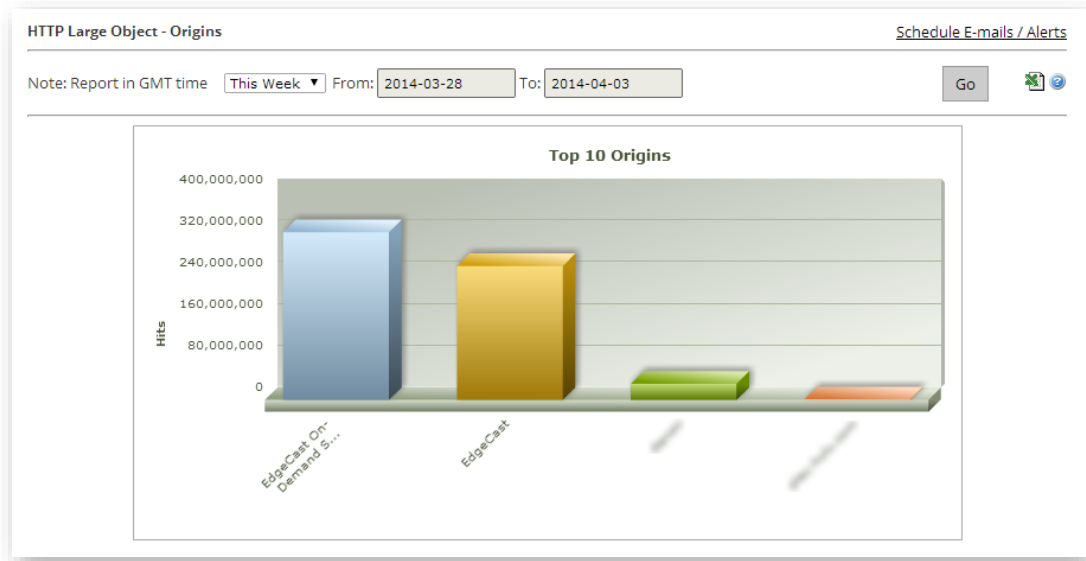
Cnames Report for the HTTP Large Platform

The data that was used to generate the chart, along with detailed bandwidth and traffic information, is displayed directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

Origins Report

The Origins report contains a graph that displays the top 10 CDN or customer origin servers from which assets were requested over a specified period of time. Statistics for the top 100 requested CDN or customer origin servers are displayed directly below this graph. Customer origin servers are identified by the name defined in the **Directory Name** option.

Note: This report uses the "EdgeCast" label to identify the CDN origin server.



Origins Report for the HTTP Large Platform


The data that was used to generate the chart, along with detailed bandwidth and traffic information, is displayed directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

Geo POPs Report

The Geo POPs report shows how much of your traffic is being routed through a particular point-of-presence (POP). The three-letter abbreviation represents a POP in our CDN network.

A list of POPs, their corresponding abbreviation, and the region that they serve is available from the CDN Help Center:

- [POP Listing](#)



Geo POPs Report for the HTTP Large Platform

The data that was used to generate the chart, along with detailed bandwidth and traffic information, is displayed directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

Clients Report

The Clients report contains a graph that displays the top 10 clients that requested assets over a specified period of time. For the purposes of this report, all requests that originate from the same IP address are considered to be from the same client. Statistics for the top 100 clients are displayed directly below this graph. This report is useful for determining download activity patterns for your top clients.

Note: Each client is identified by its IP address.



Clients Report for the HTTP Large Platform

The data that was used to generate the chart, along with detailed bandwidth and traffic information, is displayed directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

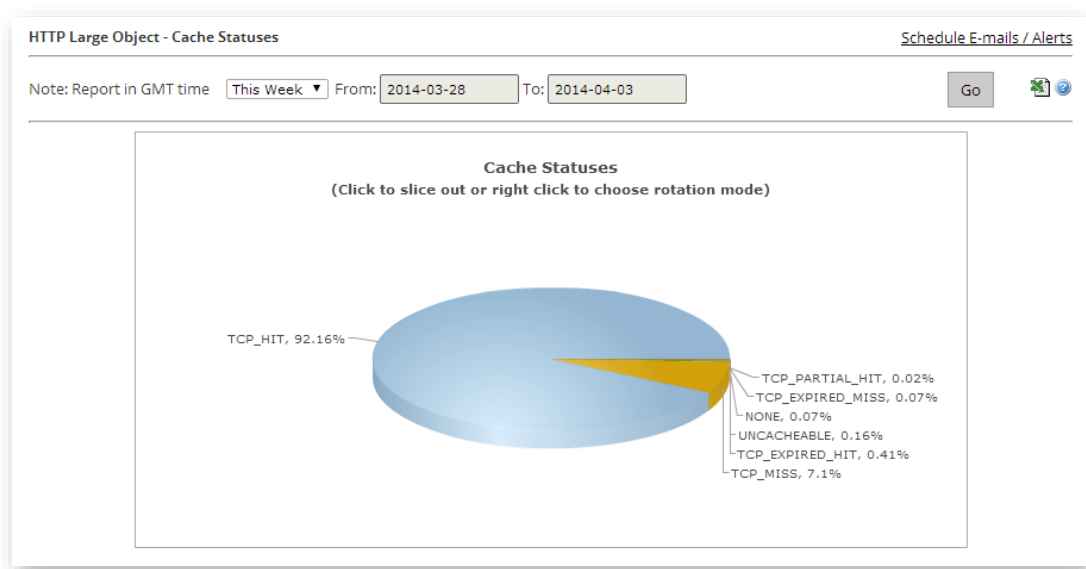
Cache Statuses Report

The Cache Statuses report gives a detailed breakdown of cache behavior, which may reveal approaches for improving the overall end-user experience. Since the fastest performance comes from cache hits, you can optimize data delivery speeds by minimizing cache misses and expired cache hits.

When reviewing this report, keep in mind that the fastest performance can be achieved when the asset is already cached on an edge server. As a result, you would like to see the majority of your cache statuses be TCP_HIT. You can improve the percentage of cache hits by performing the following:

- Configure your origin server to avoid assigning "no-cache" response headers except where strictly needed.
- Avoid expired cache hits by extending the max-age as long as possible. This will minimize the number of requests that must go to the origin server.
- Avoid query-string caching except where strictly needed.
- Avoid uncacheable response codes.

There are several possible results when checking whether an asset's content is fresh. This is known as cache status. A description for each cache status is provided in **Appendix C: Cache Statuses**.




Cache Statuses Report for the HTTP Large Platform

The data that was used to generate the chart, along with detailed bandwidth and traffic information, is displayed directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

NONE Details Report

The NONE Details report contains a graph that displays the top 10 URLs for assets for which cache content freshness was not checked over a specified period of time. Statistics for the top 100 URLs for these types of assets are displayed directly below this graph.

A cache status of NONE is reported if a cache content freshness check is not performed in response to a request for an asset. This check is skipped when Token-Based Authentication denies a request or when an HTTP request method that bypasses cache (e.g., PUT, DELETE, etc) is used.




NONE Details Report for the HTTP Large Platform

The data that was used to generate the chart, along with detailed bandwidth and traffic information, is displayed directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

CONFIG_NOCACHE Details Report

The CONFIG_NOCACHE Details report contains a graph that displays the top 10 URLs for assets that were not cached due to the customer's CDN configuration. These types of assets were served directly from the origin server. Statistics for the top 100 URLs for these types of assets are displayed directly below this graph.

A CONFIG_NOCACHE status code indicates that a customer-specific configuration on our edge servers prevented the asset from being cached. For example, an HTTP Rules Engine rule can prevent an asset from being cached by enabling the Bypass Cache feature for qualifying requests.




CONFIG_NOCACHE Details Report for the HTTP Small Platform

The data that was used to generate the chart, along with detailed bandwidth and traffic information, is displayed directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

UNCACHEABLE Details Report

The UNCACHEABLE Details report contains a graph that displays the top 10 URLs for assets that could not be cached due to request header data. Statistics for the top 100 URLs for these types of assets are displayed directly below this graph.

An UNCACHEABLE status code is reported when an asset's Cache-Control and Expires headers indicate that it should not be cached on a POP or by the HTTP client. As a result, the asset was served from the origin server.



UNCACHEABLE Details Report for the HTTP Small Platform


The data that was used to generate the chart, along with detailed bandwidth and traffic information, is displayed directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

TCP_HIT Details Report

The TCP_HIT Details report contains a graph that displays the top 10 URLs for assets that are served immediately from cache. Statistics for the top 100 URLs for these types of assets are displayed directly below this graph.

An asset has a cache status of TCP_HIT when it is served immediately from the POP to the client. An asset is immediately served from a POP when it is cached on the POP closest to the client and it has a valid time to live (TTL).

Note: TTL is determined by the Cache-Control: s-maxage, Cache-Control: max-age, and Expires headers.




TCP_HIT Details Report for the HTTP Large Platform

The data that was used to generate the chart, along with detailed bandwidth and traffic information, is displayed directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

TCP_MISS Details Report

The TCP_MISS Details report contains a graph that displays the top 10 URLs for assets that have a cache status of TCP_MISS. Statistics for the top 100 URLs for these types of assets are displayed directly below this graph.

An asset has a cache status of TCP_MISS when it could not be found on the POP closest to the client. As a result, the asset had to be requested from either an origin server or an origin shield server. If the origin server or the origin shield server returns an asset, it will be served to the client and cached on both the client and the edge server. Otherwise, a non-200 status code (e.g., 403 Forbidden, 404 Not Found, etc.) will be returned.



TCP_MISS Details Report for the HTTP Large Platform

The data that was used to generate the chart, along with detailed bandwidth and traffic information, is displayed directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

TCP_EXPIRED_HIT Details Report

The TCP_EXPIRED_HIT Details report contains a graph that displays the top 10 URLs for stale assets that were served directly from the POP. Statistics for the top 100 URLs for these types of assets are displayed directly below this graph.

An asset has a cache status of TCP_EXPIRED_HIT when an expired cached asset is served from the POP to the client. An expired cached asset (i.e., an asset whose max-age has been exceeded) is served when the origin server indicates that a newer version of the asset does not exist.



TCP_EXPIRED_HIT Details Report for the HTTP Large Platform

The data that was used to generate the chart, along with detailed bandwidth and traffic information, is displayed directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

TCP_EXPIRED_MISS Details Report

The TCP_EXPIRED_MISS Details report contains a graph that displays the top 10 URLs for stale assets for which a new version had to be retrieved from the origin server. Statistics for the top 100 URLs for these types of assets are displayed directly below this graph.

An asset has a cache status of TCP_EXPIRED_MISS when a newer version of an expired cached asset is served from the POP to the client. This occurs when the TTL for a cached asset has expired (e.g., expired max-age) and the origin server returns a newer version of that asset. This new version of the asset will be served to the client instead of the cached version. Additionally, it will be cached on the edge server and the client.

Note: An example of an expired cached asset is an asset that contains a max-age that has been exceeded.



TCP_EXPIRED_MISS Details Report for the HTTP Large Platform

The data that was used to generate the chart, along with detailed bandwidth and traffic information, is displayed directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

TCP_CLIENT_REFRESH_MISS Details Report

The TCP_CLIENT_REFRESH_MISS Details report contains a bar chart that displays the top 10 URLs for assets were retrieved from an origin server due to a no-cache request from the client. Statistics for the top 100 URLs for these types of requests are displayed directly below this chart.

An asset has a cache status of TCP_CLIENT_REFRESH_MISS when an HTTP client (e.g., browser) forces an edge server to retrieve a new version of a stale asset from the origin server. By default, our servers prevent an HTTP client from forcing our edge servers to retrieve a new version of the asset from the origin server. However, this behavior can be overridden through the use of the HTTP Rules Engine feature called "Honor No-Cache Request."

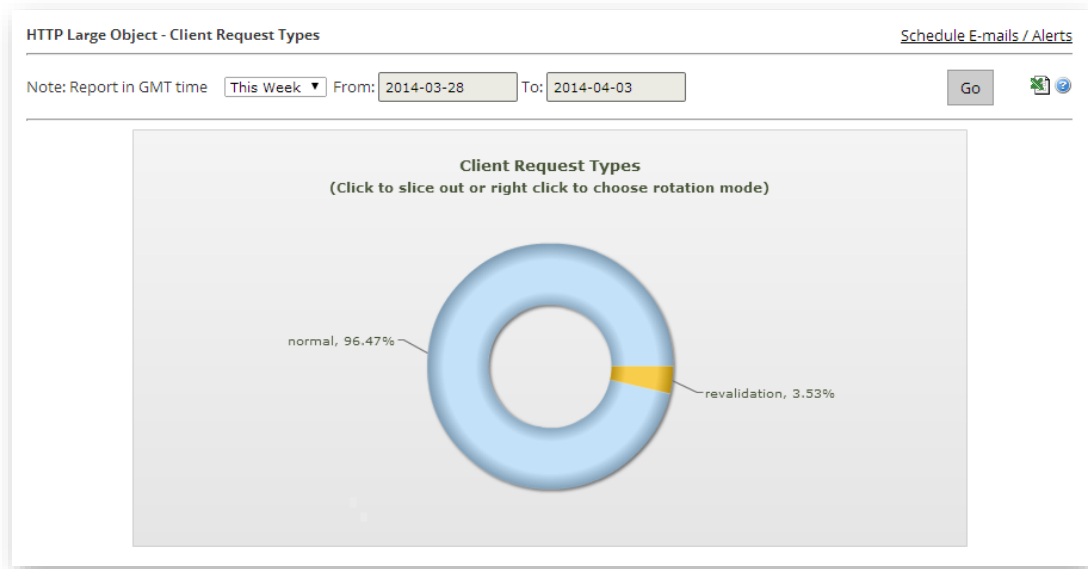
The data that was used to generate the chart, along with detailed bandwidth and traffic information, is displayed directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

Client Request Types Report

The Client Request Types report indicates the type of requests that were made by HTTP clients (e.g., browsers). This report includes a donut chart that provides a sense as to how requests are being handled. Bandwidth and traffic information for each request type is displayed below the chart.

Request Type	Description
Normal	<p>This request type is reported when all of the following conditions are met:</p> <ul style="list-style-type: none">• The requested asset must be successfully delivered to a client.• The request cannot contain a revalidation or refresh header (as defined below). <p>This type of request will return a 200 OK status code to the client.</p>
Refresh	<p>This request type is reported when all of the following conditions are met:</p> <ul style="list-style-type: none">• The requested asset must be successfully delivered to a client.• The request must contain one of the following header values:<ul style="list-style-type: none">▪ Cache-Control: no-cache▪ Pragma: no-cache <p>The use of one of the above directives forced the asset to be refreshed from the origin server. This type of request will return a 200 OK status code to the client.</p>

Request Type	Description
Revalidation	<p>This request type is reported when all of the following conditions are met:</p> <ul style="list-style-type: none"> • The requested asset must be successfully delivered to a client. • The request must contain one of the following header values: <ul style="list-style-type: none"> ▪ If-Modified-Since ▪ If-None-Match <p>The status code that is returned by this request type is determined by whether the HTTP client's cached version of the requested asset matches the one stored on an edge server. If the HTTP client contains a cached version of the asset and it matches the version on an edge server, then a 304 Not Modified response code is returned to the client. Otherwise, a new version of the asset is retrieved from the origin server and a 200 OK response code is returned to the client.</p>
Revalidation + Refresh	<p>This request type is reported when all of the following conditions are met:</p> <ul style="list-style-type: none"> • The requested asset must be successfully delivered to a client. • The request must contain both revalidation and refresh headers (as defined above). <p>The status code that is returned by this request type is determined by whether the HTTP client's cached version of the requested asset matches the one stored on the origin server. If the origin server contains the same version of the asset, then a 304 Not Modified response code is returned to the client. Otherwise, a new version of the asset is retrieved from the origin server and a 200 OK response code is returned to the client.</p>




Client Request Types Report for the HTTP Large Platform

The data that was used to generate the chart, along with detailed bandwidth and traffic information, is displayed directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

User Agents Report

The User Agent report contains a bar graph displaying the top 10 user agents to request your content through our CDN. Typically, a user agent is a web browser, media player, or a mobile phone browser. Statistics for the top 100 user agents are displayed directly below this chart.



User Agents Report for the HTTP Large Platform


The data that was used to generate the chart, along with detailed bandwidth and traffic information, is displayed directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

Referrers Report

The Referrers report contains a bar graph displaying the top 10 referrers to content accessed through our CDN. Typically, a referrer is the URL of the web page or resource that links to your content. Detailed information is provided below the graph for the top 100 referrers.

Note: By default, referrer data is not stored for the HTTP Small and ADN platforms. As a result, this report is only available for the HTTP Large platform.

The referrer that linked to your content is typically indicated as a field in the HTTP header for the requested asset. A dash (-) referrer indicates that the content was hit directly, the referrer was stripped out by the user agent, or the referrer was not passed through the HTTP header.



Referrers Report for the HTTP Large Platform

The data that was used to generate the chart, along with detailed bandwidth and traffic information, is displayed directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

Compression Types Report

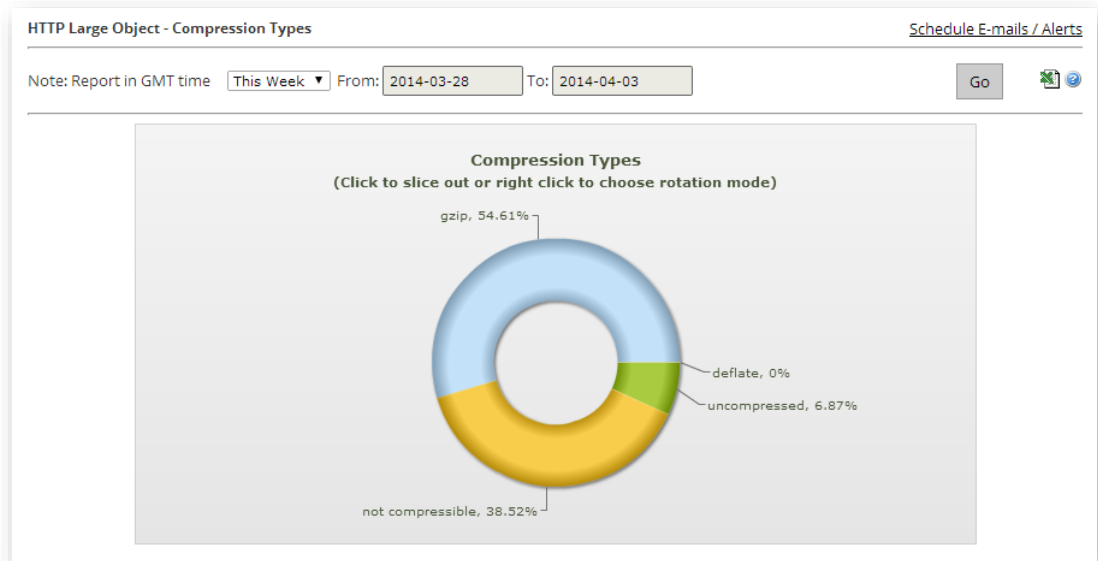
The Compression Types report contains a donut chart that breaks down requested assets by whether they were compressed by our edge servers. The percentage of compressed assets is broken down by the type of compression used. Detailed information is provided below the graph for each compression type and status.

Before an asset will be compressed by our CDN, the following conditions must be met:

1. The platform (i.e., HTTP Large, HTTP Small, or ADN) through which the asset was requested must be configured to allow compression for the asset's Internet media type (i.e., mime-type).
2. The HTTP client must send an Accept-Encoding header with at least one acceptable compression type (i.e., gzip, deflate, or bzip2).
3. The edge server handling the request must already have a cached version of the requested asset.
4. The requested asset must be smaller than 1 MB, which is the maximum allowed size for compression.

The available compression types are described below.

Compression Type	Description
bzip2	The requested asset met the above conditions and was compressed using the bzip2 format.
deflate	The requested asset met the above conditions and was compressed using the deflate format.
gzip	The requested asset met the above conditions and was compressed using the gzip format.
not compressible	The Internet media type for the requested asset was not enabled for compression on the HTTP Large platform
uncompressed	The requested asset met condition #1, but did not meet one of the other requirements. As a result, the asset was not compressed.




Compression Types Report for the HTTP Large Platform

The data that was used to generate the chart, along with detailed bandwidth and traffic information, is displayed directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

File Types Report

The File Types report contains a bar graph that displays the top 10 file types that have been requested through our CDN for your account. For the purposes of this report, a file type is defined by the asset's file name extension and Internet media type (e.g., .html [text/html], .htm [text/html], .aspx [text/html], etc.). Detailed information is provided below the graph for the top 100 file types.

Note: Keep in mind that multiple Internet media types can be associated with a single file name extension. As a result, each unique combination of file name extension and Internet media type is considered a different file type.



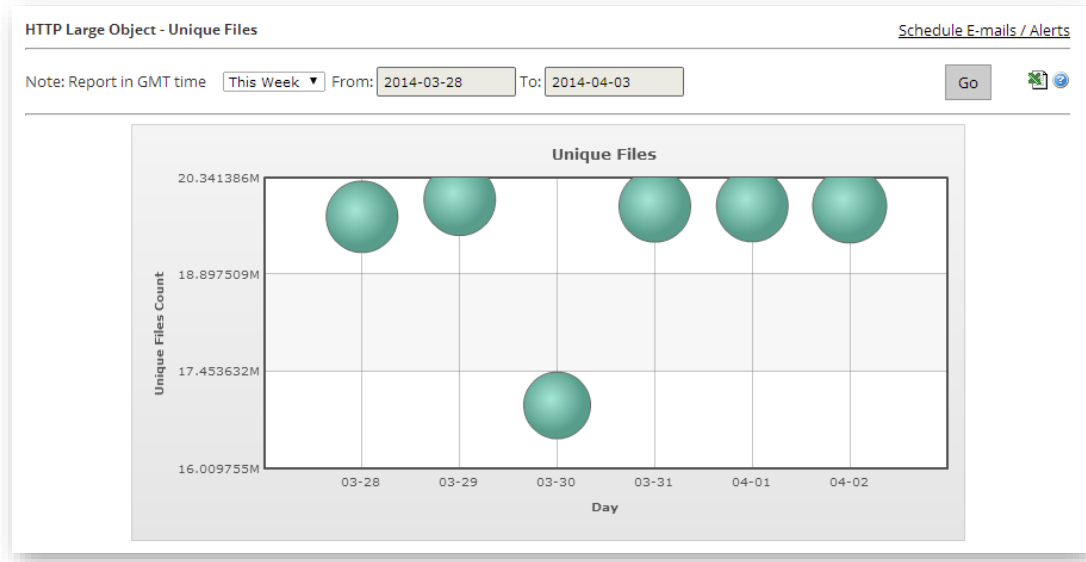
File Types Report for the HTTP Large Platform

The data that was used to generate the chart, along with detailed bandwidth and traffic information, is displayed directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

Unique Files Report

The Unique Files report contains a graph that plots the total number of unique assets that were requested on a particular day over a specified period of time.

Tip: Hover the cursor over a sphere to view the date (in Unix time), the total number of unique files that were requested on that date, and the total number of bytes that were transferred for those unique files on that date.



Unique Files Report for the HTTP Large Platform

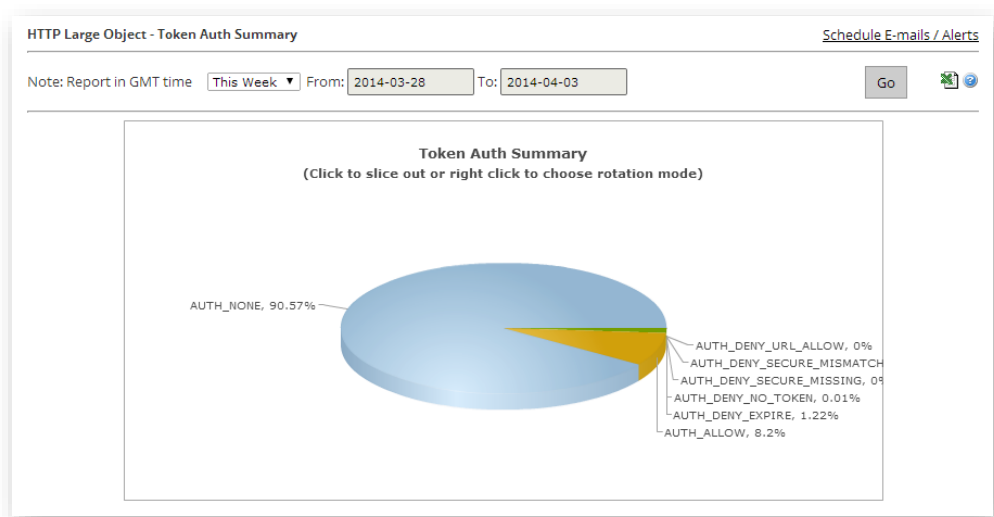
The data that was used to generate the chart is displayed directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

Token Auth Summary Report

The Token Auth Summary report contains a pie chart that provides a quick overview on whether requested assets were protected by Token-Based Authentication. Protected assets are displayed in the chart according to the results of their attempted authentication. The available Token-Based Authentication status codes are described below.

Status Code	Description
AUTH_ALLOW	This status code is used to keep track of all requested assets that were successfully authenticated through Token-Based Authentication.
AUTH_DENY_CLIENTIP	This status code is used to keep track of all requested assets for which access was denied as a result of the ec_clientip parameter.
AUTH_DENY_COUNTRY	This status code is used to keep track of all requested assets for which access was denied as a result of the ec_country parameter.
AUTH_DENY_COUNTRY_ALLOW	This status code is used to keep track of all requested assets for which access was denied as a result of the ec_country_allow parameter.
AUTH_DENY_COUNTRY_DENY	This status code is used to keep track of all requested assets for which access was denied as a result of the ec_country_deny parameter.
AUTH_DENY_EXPIRE	This status code is used to keep track of all requested assets for which access was denied as a result of the ec_expire parameter.
AUTH_DENY_NO_KEY	This status code is used to keep track of all requested assets for which access was denied because a primary or backup key was not specified for Token-Based Authentication on the HTTP Large platform.
AUTH_DENY_NO_TOKEN	This status code is used to keep track of all requested assets for which access was denied because a token was not specified as a query string in the URL.
AUTH_DENY_PROTO_ALLOW	This status code is used to keep track of all requested assets for which access was denied as a result of the ec_proto_allow parameter.
AUTH_DENY_PROTO_DENY	This status code is used to keep track of all requested assets for which access was denied as a result of the ec_proto_deny parameter.

Status Code	Description
AUTH_DENY_REF	This status code is used to keep track of all requested assets for which access was denied as a result of the ec_ref parameter.
AUTH_DENY_REF_ALLOW	This status code is used to keep track of all requested assets for which access was denied as a result of the ec_ref_allow parameter.
AUTH_DENY_REF_DENY	This status code is used to keep track of all requested assets for which access was denied as a result of the ec_ref_deny parameter.
AUTH_DENY_SECURE_MISMATCH	This status code is used to keep track of all requested assets for which access was denied because the value in ec_secure didn't match the length of the actual decrypted token string.
AUTH_DENY_SECURE_MISSING	This status code is used to keep track of all requested assets for which access was denied because ec_secure was not found in the decrypted token string
AUTH_DENY_URL_ALLOW	This status code is used to keep track of all requested assets for which access was denied as a result of the ec_url_allow parameter.
AUTH_DENY_URL_ALLOW_EMPTY	This status code is used to keep track of all requested assets for which access was denied due to an empty value being specified for the ec_url_allow parameter.
AUTH_NONE	This status code is used to keep track of all requested assets that were not protected by Token-Based Authentication.

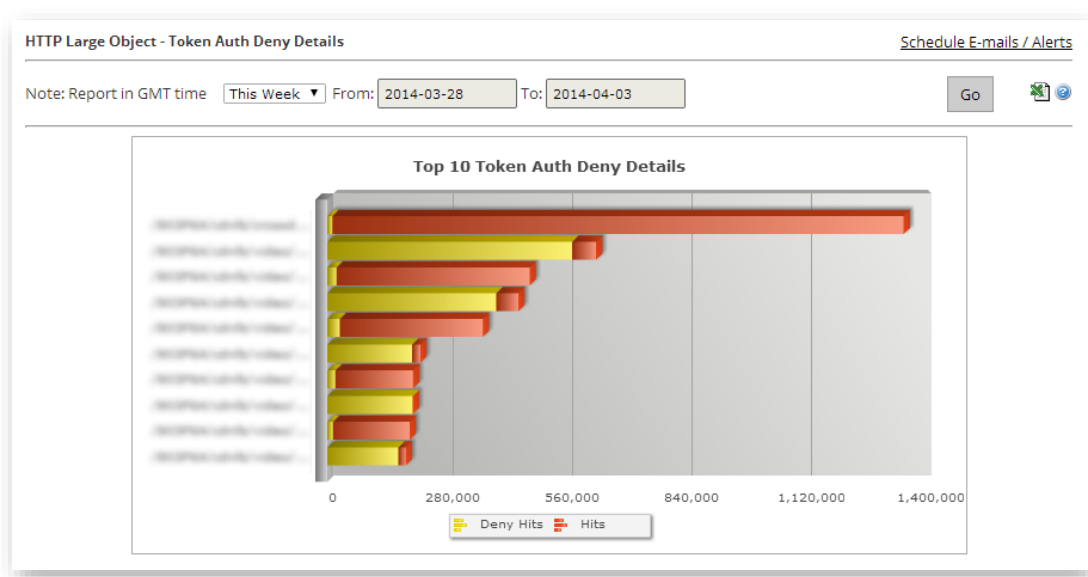


Token Auth Summary Report for the HTTP Large Platform

Additional information is provided directly below the pie chart. Those fields are described in **Appendix A: Edge Performance Analytic Fields**.

Token Auth Deny Details Report

The Token Auth Deny Details report contains a bar graph that allows you to view the top 10 requests that were denied due to Token-Based Authentication. For each denied request, you can get a quick sense of how many requests were denied for a particular asset by comparing the color-coded portions of the bar corresponding to it. The yellow portion of the bar indicates all of the requests that were denied for that asset. The remaining red portion represents the requests that were authorized for that asset. You can view the total hits, the number of denied requests, and the percentage of denied requests below the graph.




Token Auth Deny Details Report for the HTTP Large Platform

Additional information is provided directly below the pie chart. Those fields are described in **Appendix A: Edge Performance Analytic Fields**.

HTTP Response Codes

The HTTP Response Codes report provides a breakdown of the HTTP status codes (e.g., 200 OK, 403 Forbidden, 404 Not Found, etc.) that were delivered to your HTTP clients by our edge servers. A pie chart allows you to quickly assess how your assets were served. Detailed statistical data is provided for each response code below the graph.



HTTP Response Codes Report for the HTTP Large Platform

The data that was used to generate the chart is displayed directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

400 Errors

The 400 Errors report contains a bar graph that allows you to view the top 10 requests that resulted in a 400 Bad Request response code. A 400 Bad Request response code occurs when your web server refuses a malformed request. For each asset in the bar graph, you can get a quick sense of how many requests resulted in a 400 Bad Request response code by comparing the color-coded portions of the bar corresponding to it. The yellow portion of the bar indicates all of the requests that returned a 400 Bad Request response code for that asset. The remaining red portion represents the requests that returned a different response code for that asset.

This report displays the data corresponding to the graph directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

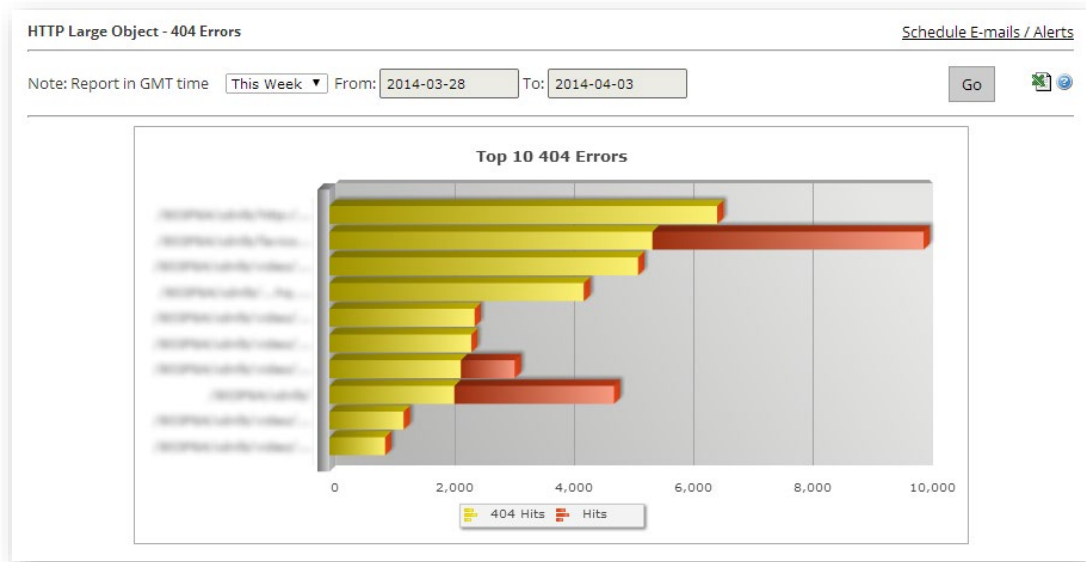
401 Errors

The 401 Errors report contains a bar graph that allows you to view the top 10 requests that resulted in a 401 Unauthorized response code. A 401 Unauthorized response code occurs when the request does not contain the proper authentication credentials for content requested from your web server. For each asset in the bar graph, you can get a quick sense of how many requests resulted in a 401 Unauthorized response code by comparing the color-coded portions of the bar corresponding to it. The yellow portion of the bar indicates all of the requests that returned a 401 Unauthorized response code for that asset. The remaining red portion represents the requests that returned a different response code for that asset.

This report displays the data corresponding to the graph directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

404 Errors

The 404 Errors report contains a bar graph that allows you to view the top 10 requests that resulted in a 404 Not Found response code. A 404 Not Found response code occurs when the requested asset does not exist. For each asset in the bar graph, you can get a quick sense of how many requests resulted in a 404 Not Found response code by comparing the color-coded portions of the bar corresponding to it. The yellow portion of the bar indicates all of the requests that returned a 404 Not Found response code for that asset. The remaining red portion represents the requests that returned a different response code for that asset.



404 Errors Report for the HTTP Large Platform


The data that was used to generate the graph is displayed directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

403 Errors

The 403 Errors report contains a bar graph that allows you to view the top 10 requests that resulted in a 403 Forbidden response code. A 403 Forbidden response code occurs when a request is denied by a customer origin server or an edge server on our POP.

Note: An edge server will deny a request when the requested asset cannot be validated through Token-Based Authentication.

For each asset in the bar graph, you can get a quick sense of how many requests resulted in a 403 Forbidden response code by comparing the color-coded portions of the bar corresponding to it. The yellow portion of the bar indicates all of the requests that returned a 403 Forbidden response code for that asset. The remaining red portion represents the requests that returned a different response code for that asset.



403 Errors Report for the HTTP Large Platform

The data that was used to generate the graph is displayed directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

429 Errors

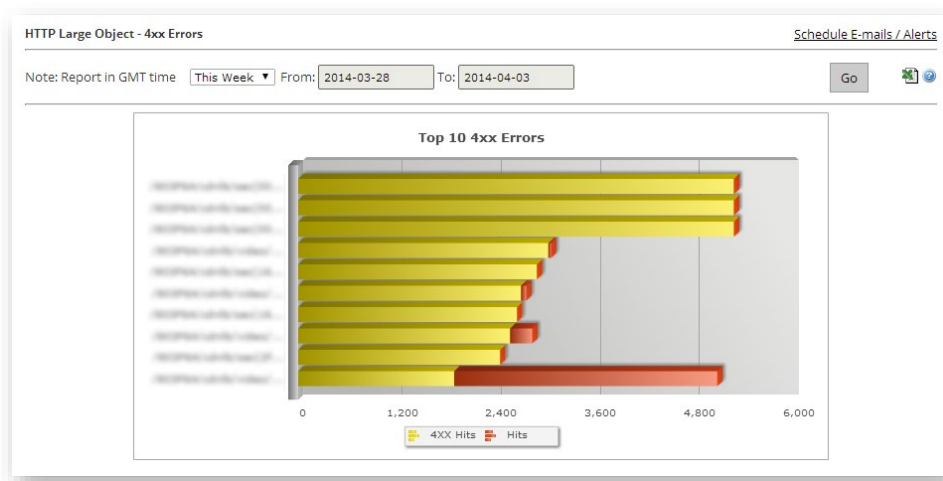
The 429 Errors report contains a bar graph that allows you to view the top 10 requests that resulted in a 429 Too Many Requests response code. A 429 Too Many Requests response code occurs when the number of requests submitted by the client exceeds your web server's rate limiting policy. For each asset in the bar graph, you can get a quick sense of how many requests resulted in a 429 Too Many Requests response code by comparing the color-coded portions of the bar corresponding to it. The yellow portion of the bar indicates all of the requests that returned a 429 Too Many Requests response code for that asset. The remaining red portion represents the requests that returned a different response code for that asset.

This report displays the data corresponding to the graph directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

4xx Errors

The 4xx Errors report contains a bar graph that allows you to view the top 10 requests that resulted in a response code in the 400 range. Typically, a 4xx response code occurs when a request is denied as a result of a client error.

For each asset in the bar graph, you can get a quick sense of how many requests resulted in a 4xx response code by comparing the color-coded portions of the bar corresponding to it. The yellow portion of the bar indicates all of the requests that returned a 4xx response code for that asset. The remaining red portion represents the requests that returned a different response code for that asset.



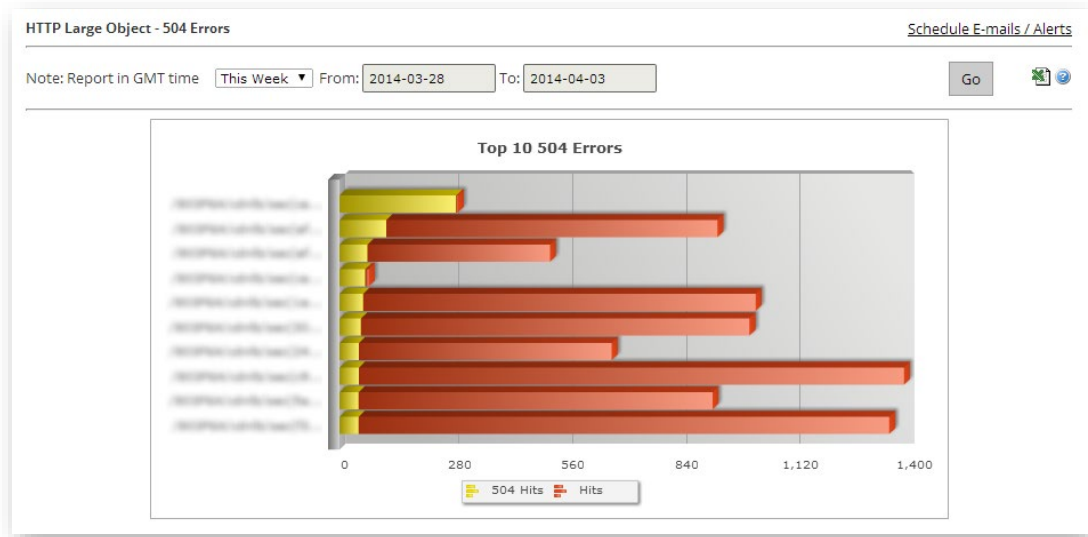
4xx Errors Report for the HTTP Large Platform

The data that was used to generate the graph is displayed directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

504 Errors

The 504 Errors report contains a bar graph that allows you to view the top 10 requests that resulted in a 504 Gateway Timeout response code. A 504 Gateway Timeout response code occurs when a timeout occurs when an HTTP proxy is trying to communicate with another server. In the case of our CDN, a 504 Gateway Timeout response code typically occurs when an edge server is unable to establish communication with a customer origin server.

For each asset in the bar graph, you can get a quick sense of how many requests resulted in a 504 Gateway Timeout response code by comparing the color-coded portions of the bar corresponding to it. The yellow portion of the bar indicates all of the requests that returned a 504 Gateway Timeout response code for that asset. The remaining red portion represents the requests that returned a different response code for that asset.



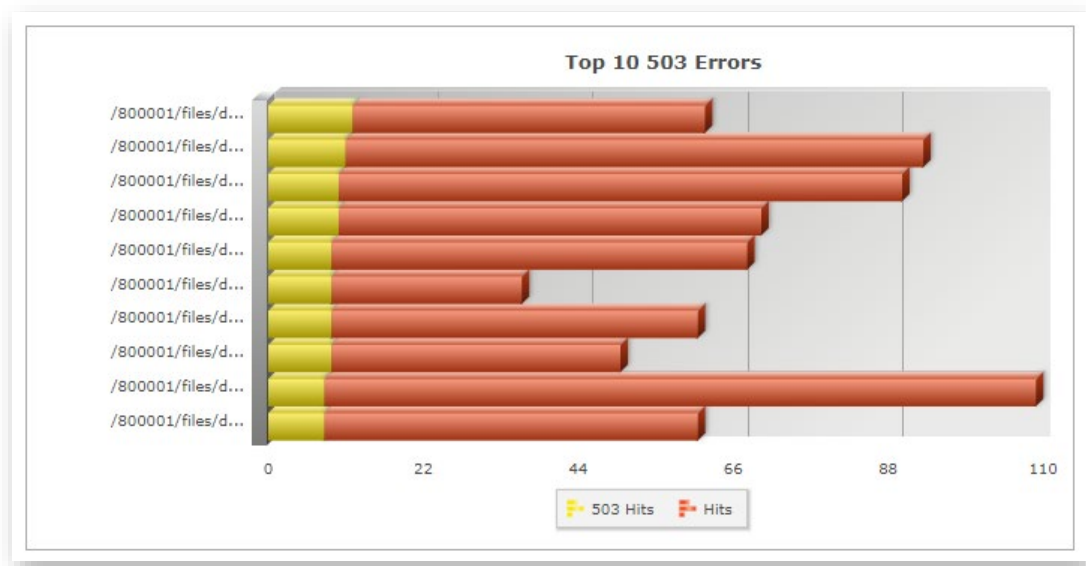
504 Errors Report for the HTTP Large Platform

The data that was used to generate the graph is displayed directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

503 Errors

The 503 Errors report contains a bar graph that allows you to view the top 10 requests that resulted in a 503 Service Unavailable response code. A 503 Service Unavailable response code occurs when the origin server is unable to provide a response. In the case of our CDN, a 503 Service Unavailable response code typically occurs when a web server associated with a customer origin configuration is unavailable either due to maintenance or because it is overloaded.

For each asset in the bar graph, you can get a quick sense of how many requests resulted in a 503 Service Unavailable response code by comparing the color-coded portions of the bar corresponding to it. The yellow portion of the bar indicates all of the requests that returned a 503 Service Unavailable response code for that asset. The remaining red portion represents the requests that returned a different response code for that asset.



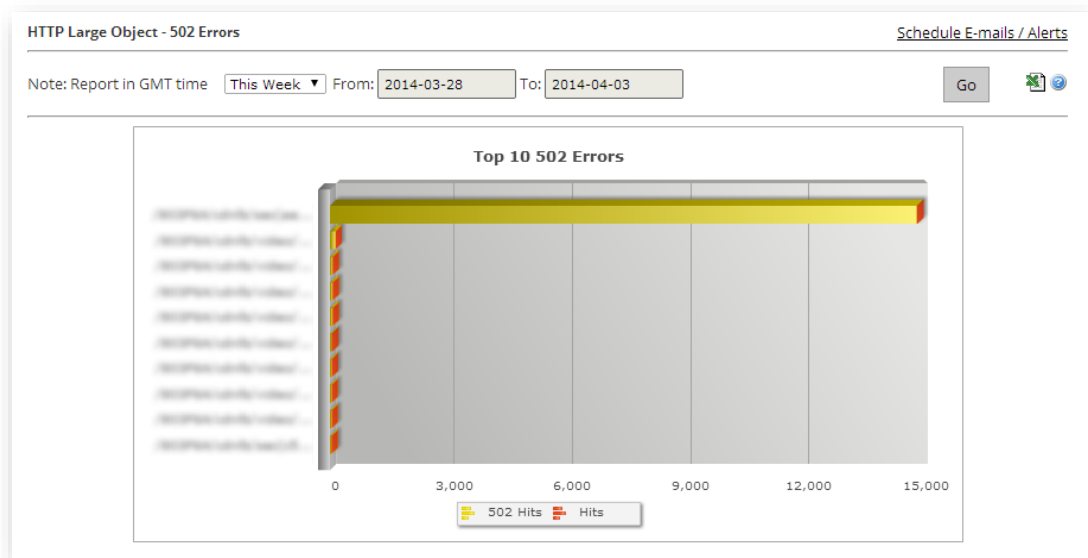
503 Errors Report for the HTTP Large Platform

The data that was used to generate the graph is displayed directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

502 Errors

The 502 Errors report contains a bar graph that allows you to view the top 10 requests that resulted in a 502 Bad Gateway response code. A 502 Bad Gateway response code occurs when an HTTP protocol failure occurs between a server and an HTTP proxy. In the case of our CDN, a 502 Bad Gateway response code typically occurs when a customer origin server returns an invalid response to an edge server. A response is invalid if it cannot be parsed or if it is incomplete.

For each asset in the bar graph, you can get a quick sense of how many requests resulted in a 502 Bad Gateway response code by comparing the color-coded portions of the bar corresponding to it. The yellow portion of the bar indicates all of the requests that returned a 502 Bad Gateway response code for that asset. The remaining red portion represents the requests that returned a different response code for that asset.



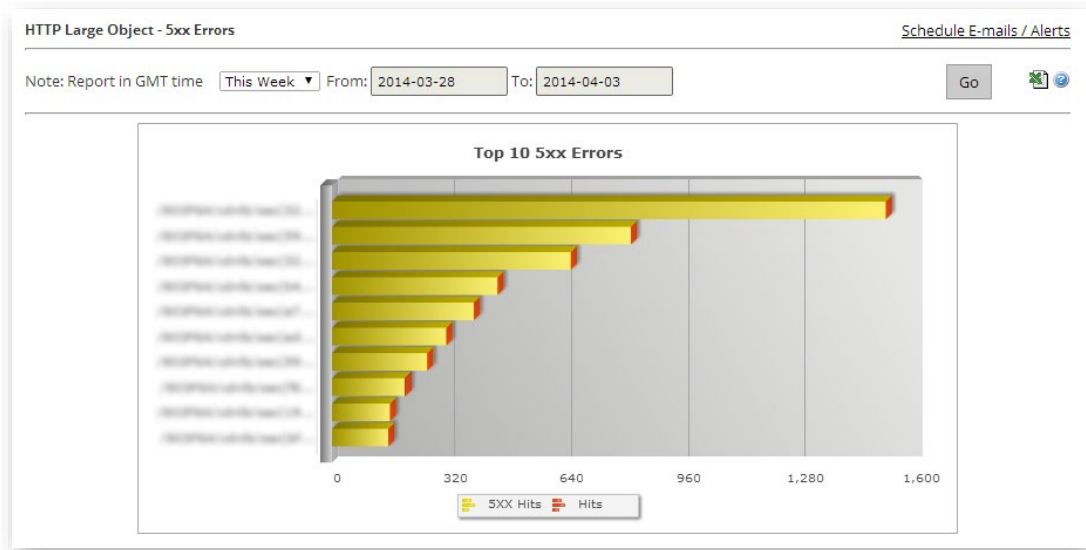
502 Errors Report for the HTTP Large Platform

The data that was used to generate the graph is displayed directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

5xx Errors

The 5xx Errors report contains a bar graph that allows you to view the top 10 requests that resulted in a response code in the 500 range. Typically, a 5xx response code occurs when a request is denied as a result of a server error. In CDN usage, the vast majority of response codes in the 500 range are generated from a customer origin server and proxied through an edge server to the HTTP client.

For each asset in the bar graph, you can get a quick sense of how many requests resulted in a 5xx response code by comparing the color-coded portions of the bar corresponding to it. The yellow portion of the bar indicates all of the requests that returned a 5xx response code for that asset. The remaining red portion represents the requests that returned a different response code for that asset.



5xx Errors Report for the HTTP Large Platform

The data that was used to generate the graph is displayed directly below it. These fields are described in **Appendix A: Edge Performance Analytic Fields**.

Real-Time Alerts

Overview

A real-time alert is designed to provide you with real-time notification as to current trends in the network traffic for your account. Before you can start receiving notifications, you will need to define one or more real-time alerts. Each alert consists of the following elements:

- Alert Criteria
- Alert Notification Definition

Note: In order to configure real-time alerts, you will need the Reports: Real-Time Stats privilege. However, any user with a valid e-mail address can receive a real-time alert notification.

Alert Criteria

Alert criteria consist of a platform and the condition that must be met before a notification will be sent out. The specified condition will only be monitored on the selected platform. For example, you can define an alert to notify you when more than a certain number of 5xx status codes are being returned to your customers every second on the HTTP Large platform.

Alert Notification Definition

An alert also defines the method and the content of the notification that will be sent out. You can be notified by e-mail or by posting an HTTP request. Both types of notifications can be customized.

Alert Criteria

The most important aspect of a real-time alert configuration is the definition of what will trigger an alert. It determines the type of CDN conditions that must exist before a notification will be sent out. The three options that play a vital role in determining when a real-time alert will be triggered are:

- Platform
- Expression
- Notify On

Tip: Detailed information on the available alert criteria options can be found in the **User Interface** section below.

Platform

Only the platform associated with an alert will be monitored for the specified condition. If you would like to monitor several platforms for the same condition, then you will need to create an alert for each desired platform.

Expression

An expression determines the type of CDN activity that will trigger an alert. It consists of a metric, a mathematical operator, and the trigger value. The metric defines what type of condition will be monitored. The mathematical operator (e.g., >, <, =, etc.) establishes the relationship between the metric and the trigger value. The trigger value determines the threshold value that must be met before a notification will be sent out.

Notify On

The **Notify on** option determines when a notification will be sent after the specified CDN condition is detected. This option determines whether a notification will be sent when it is first detected, each time it is detected, and/or when it is no longer detected.

Notifications

There are two types of actions that can take place when a specified condition has been met, which are e-mail notifications or posting data to an external server. E-mail notifications allow you to stay on top of CDN account activity without having to monitor real-time statistics for the desired platform. HTTP (POST) notifications allow you to leverage a web server to transform data provided by our network monitoring system.

E-mail Notifications

An alert can be configured to send an e-mail whenever a predefined condition has been met. For example, if you configured an alert to set off when your bandwidth exceeds 10000 Mbps, then a system e-mail will be sent as soon as your bandwidth reaches 10001 Mbps.

An e-mail notification triggered by an alert consists of a standard e-mail format. It consists of a sender, recipients, a subject line, and the body. Each of these elements is examined below.

Sender

The sender of this type of e-mail is a predefined system e-mail address. This predefined e-mail address cannot be altered.

Tip: If you are not receiving alerts, then your e-mail provider may have tagged e-mails from this sender as junk mail. In such a case, you will need to indicate that e-mails from this sender are safe and should be sent to your Inbox.

Recipients

The recipients of an e-mail notification are defined on a per alert basis. When defining the recipients of an alert's e-mail notification, you can specify a single or multiple e-mail addresses. If you decide to specify more than one recipient, make sure to separate each one with a comma (e.g., jsmith@hotmail.com, jsmith@gmail.com, jsmith@yahoo.com). The blank space after the comma is optional.

Subject Line

The subject line of an e-mail notification should provide a concise summary of the type of alert that has been triggered. For example, you could set the subject line to read "Warning: Excessive Occurrences of 404 Not Found Status Codes," when the number of 404 occurrences per second warrants your attention.

In addition to providing static information that identifies an alert, you can include dynamic information that will be replaced with information that is specific to the conditions that triggered the alert. This can be accomplished through the use of keywords. Keywords allow you to tailor the subject line to provide more information at a glance. For example, the above subject line could be reworked to read the following, "Warning: [CurrentValue] Occurrences per Second of the 404 Not Found Status Code Were Detected." The keyword [CurrentValue] would be replaced with the value that triggered the alert. An e-mail triggered by this alert could contain the following subject line, "Warning: 10 Occurrences per Second of the 404 Not Found Status Code Were Detected."

Tip: For more information on keywords, please refer to the **Notification Keywords** section below.

Body

The body of an e-mail notification allows you to provide more detailed information about the alert that was triggered. Similar to the subject line, the body can consist of static and dynamic text. You can specify dynamic text through the use of keywords.

HTTP (Post) Notifications

An HTTP POST notification provides the means through which you can post data about an alert on a server outside of our network. An HTTP POST notification consists of a URL, request headers, and a request body. Keywords can be used to define the request body. A keyword is replaced with dynamic data when the POST request is sent.

Reminder: For more information on keywords, please refer to the **Notification Keywords** section below.

Since you are defining the URL to which a POST request will be sent, you can define the request handler that will accept it. This allows you to take advantage of your preferred scripting language to handle the POST requests sent by our network monitoring system. The possibilities

for what your script can do with this information are endless. For example, your script could perform one or more of the following actions:

- Update an existing network monitoring system via an API.
- Populate a database with detailed information about the alert. One use for such a database is to generate an alert history for your account, which you can analyze for trends or cross-reference with our existing reporting capabilities.
- Create an RSS feed that you can make available to interested customers.
- Send an e-mail notification to your customers.

Important: All of the above suggestions require programming knowledge and access to a web server.

Important: The creation, modification, or support of custom scripts or web pages is not a service that we provide.

Tip: You can test an HTTP POST notification by clicking on **Test Notification**. The status code returned by the server will be indicated directly below this option.

Notification Keywords

Keywords are system-defined words that are dynamically replaced when an alert notification is sent out. This allows a notification to contain information that is specific to the alert that was triggered. A brief description is provided for each available keyword below.

Keyword	Description
[AlertDuration]	Indicates the length of time (in minutes) that the condition that triggered the alert has been detected.
[CurrentValue]	Indicates the current value of the metric that was used to trigger the alert. For example, if your alert's expression is triggered by the value associated with "Bandwidth Mbps," then this keyword will be replaced by the bandwidth value in Mbps that triggered the alert (e.g., 10001).
[Expression]	Indicates the expression that was used by the alert to determine whether a notification should be sent. An expression consists of a metric, operator, and a trigger value. The format for this keyword is: <i>Metric Operator TriggerValue</i> (e.g., <i>Bandwidth Mbps > 10000</i>).

Keyword	Description
[Interval]	Indicates the time interval (in minutes) that must pass before our network monitoring system can check for the specified trigger value again. This keyword is replaced by the integer value associated with the Interval option when the alert was triggered.
[MediaType]	Indicates the platform (e.g., HTTP Large, HTTP Small, etc.) that is being checked for the specified condition. This keyword is replaced by the value assigned to the Media Type option.
[Metric]	Indicates the metric (e.g., Bandwidth Mbps, Total Connections, Cache Status: Total Hits per second, etc.) that was being monitored for a specified value. This keyword is replaced by the name of the metric associated with the alert when it was triggered. For a list of available metrics, please refer to the Appendix B: Monitoring Criteria (Metrics) .
[Name]	Indicates the name assigned to the alert when it was triggered.
[NotificationCondition]	Indicates the conditions under which a notification will be sent after an alert has been triggered. This keyword is replaced by the value assigned to the Notify on option.
[Operator]	Indicates the relational operator used to establish the threshold for triggering an alert. This keyword is replaced by the value assigned to the operator portion of the Expression option.
[TriggerValue]	Indicates the threshold value for determining when an expression is true. This keyword is replaced by the value assigned to the trigger value portion of the Expression option.

Real-Time Alert Administration

The three basic administrative tasks that you can perform with real-time alerts are:

- Creation
- Modification
- Deletion

Each of these tasks is explained in this section.

Creating a Real-Time Alert

A real-time alert can be created from the **Alerts** page in the MCC. When creating an alert, you must specify a name, a platform, and an expression. Although those are the only required options, it is recommended that you configure notification options when creating an alert. If you are not ready to receive notifications, then you can simply disable your rule by making sure that the **Alert Enabled** option is cleared.

To create a real-time alert

1. Navigate to the **Alerts** page. Load this page by finding the **Analytics** menu and then selecting **Real-Time Stats**. After which, click on **Real-Time Alerts** from the side navigation bar.
2. Click **Add Alert** to display the **Alert Configurations** page.
3. Mark the **Alert Enabled** option.
4. In the **Name** option, type a label that can be used to identify this alert. This label could be a short description that describes the purpose of this alert.
5. In the **Media Type** option, select the platform that will be monitored.
6. In the **Expression** option, perform the following steps:
 - i. At the start of the expression, you will need to select the metric that will be monitored on your platform.
 - ii. In the middle of the expression, you will need to select a relational operator that determines the relationship between the metric and the trigger value.
 - iii. At the end of the expression, you will need to specify the alert's trigger value. This is the threshold value that must be met before an alert can send off a notification.

7. From the **Cname** option, determine whether this alert applies to all traffic or a specific edge CNAME.

- **All Traffic:** From the **Cname** option, select "All Cnames."
- **Edge CNAME:** From the **Cname** option, select the desired edge CNAME configuration.

Note: The **Cname** option lists all edge CNAMEs for which custom report logging has been enabled.

8. In the **Interval** option, specify how often the network monitoring system will check for the specified condition on your account. This value is specified in minutes.

9. In the **Notify on** option, select how often a notification will be sent when the condition specified in step 5 is detected. For the purposes of this option, an instance of a condition is determined by the first time it is detected until it is no longer detected.

10. Perform one of the following:

- If you would like to send an e-mail notification when the specified condition is detected, then you should select the **Notify by Email** option.
 - i. In the **To** option, specify each recipient of an e-mail notification for this alert.
 - ii. In the **Subject** option, review the default subject line and make the desired changes.
 - iii. In the **Body** option, review the default body and make the desired changes.
- If you do not wish to send e-mail notifications, clear the **Notify by Email** option and proceed to the next step.

11. Perform one of the following:


- If you would like to send an HTTP POST notification when the specified condition is detected, then you should select the **Notify by HTTP Post** option.
 - i. In the **URL** option, specify the URL to the web page that will handle the HTTP POST request. This URL should point to the application server that will receive the request.
 - ii. In the **Headers** option, define the request headers that will be sent to the web server defined in the **URL** option. Each request header should be placed on its own line. Keep in mind that the Content-Type request header determines the format for the request body defined in the **Body** option.

- iii. In the **Body** option, define the request body that will be sent to the previously defined application server when the specified network condition is detected. Make sure to use the format defined by the Content-Type request header (e.g., XML or JSON).
- iv. Optional. If you would like to send a test HTTP request to your application server, you should click **Test Notification**.
- If you do not wish to send alert notifications via HTTP POST, clear the **Notify by HTTP Post** option and proceed to the next step.


12. Click **Save** to create a new real-time alert.

Tip: An alert can be configured to send both an e-mail notification and an HTTP POST request to the desired server whenever the specified CDN condition is detected.

Modifying a Real-Time Alert

A real-time alert can be modified by clicking the edit icon () next to its name from the **Real-Time Alerts** page. The **Alert Configurations** page will display the settings associated with the selected real-time alert. You can then modify the desired settings and then click **Save** to save your changes.

Deleting a Real-Time Alert

A real-time alert can be deleted by clicking the delete icon () next to its name from the **Real-Time Alerts** page. When prompted, confirm the deletion of the real-time alert.

Finding a Real-Time Alert

A list of real-time alerts can be found on the **Alerts** page. If the list is too large to easily browse through, then you may search by the information associated with the desired real-time alert. A search can be performed from the **Alerts** page by specifying the word that you would like to search for in the **Search** option and then clicking **Search**.

Tip: If you would like to view a listing of all real-time alerts, simply clear the **Search** option and perform a blank search. A blank search will return all real-time alerts.

User Interface

This section explains the different options that can be used to define a real-time alert configuration.

Alert Configurations

A real-time alert must be defined before our network monitoring system can monitor your account. This definition involves deciding which platform will be monitored, what will be monitored, the criteria that will be used to trigger a notification, and how often our network monitoring system will check for the specified condition. A brief explanation is provided for each of these options below.

Name	Description
Alert Enabled	Determines whether the current alert will be active or inactive. If an alert has been deactivated, then our network monitoring system will not monitor your account for the specified condition. As a result, a notification will not be sent when the condition takes place.
Name	Required. Determines how this new alert will be identified. A properly named alert provides a quick way to identify its purpose.
Media Type	Determines the platform (e.g., HTTP Large or HTTP Small) that will be monitored by the criteria defined by this alert. You may choose to monitor other platforms for the same condition by creating an alert on each desired platform.
Expression	Required. Determines the network condition that will be monitored for on the specified platform. When defining an expression, you will need to set the following options: <ul style="list-style-type: none">• Metric: Identifies the type of CDN activity/condition that will be monitored on a specified platform for your account. A complete list of metrics is provided in the Appendix B: Monitoring Criteria (Metrics).• Operator: Determines the relational operator that will be used to establish whether a particular network condition holds true.<ul style="list-style-type: none">▪ > (Greater Than): Indicates that the alert criteria will be satisfied when the current value of the metric is greater than the value specified in the Trigger value option.▪ < (Less Than): Indicates that the alert criteria will be satisfied when the current value of the metric is less than the value specified in the Trigger value option.▪ = (Equal To): Indicates that the alert criteria will be satisfied when the current value of the metric is equal to the value specified in the Trigger value option.▪ >= (Greater Than or Equal To): Indicates that the alert criteria

Name	Description
	<p>will be satisfied when the current value of the metric is greater than or equal to the value specified in the Trigger value option.</p> <ul style="list-style-type: none"> ▪ <= (Less Than or Equal To): Indicates that the alert criteria will be satisfied when the current value of the metric is less than or equal to the value specified in the Trigger value option. ▪ != (Not Equal To): Indicates that the alert criteria will be satisfied when the current value of the metric is not equal to the value specified in the Trigger value option. <ul style="list-style-type: none"> • Trigger value: Determines the threshold value for determining when an expression is true.
Filters	<p>Determines whether the current alert will only be triggered for traffic corresponding to a specific edge CNAME.</p> <hr/> <p>Note: The Cname option lists all edge CNAMEs for which custom report logging has been enabled.</p> <hr/>
Interval	<p>Determines the number of minutes that must pass between each check for the specified expression. The specified value must be a whole number greater than 0.</p>
Notify On	<p>Determines the notification frequency during the time period that the specified condition is true. You can choose between the following options:</p> <ul style="list-style-type: none"> • Condition Start: Indicates that a notification will be sent when the specified condition is first detected. • Condition End: Indicates that a notification will be sent when the specified condition is no longer detected. This notification can only be triggered after our network monitoring system detected that the specified condition occurred. • Continuous: Indicates that a notification will be sent each time that the network monitoring system detects the specified condition. Keep in mind that the network monitoring system will only check once per interval for the specified condition. • Condition Start and End: Indicates that a notification will be sent the first time that the specified condition is detected and once again when the condition is no longer detected. This option provides approximate notification as to when the specified condition started and ended.

E-mail Notification

When an alert's condition has been met, then a custom e-mail can be sent according to the **Notify On** option. The properties of this e-mail, such as recipients, the subject line, and the body, are determined by the options specified in the **Notify by Email** section.

Name	Description
Notify by Email	Determines whether an e-mail notification will be sent out when an alert's expression and notification conditions are met.
From	<p>Indicates the e-mail address from which the e-mail notification will be sent. This e-mail address is set by your CDN provider.</p> <hr/> <p>Tip: You should ensure that your e-mail system handles e-mails from this e-mail address properly to avoid a situation where alerts are blocked, filtered, or otherwise mishandled.</p> <hr/>
To	<p>Determines the set of e-mail addresses to which an e-mail notification will be sent. Only valid e-mail addresses should be specified in this option. If you would like to specify more than one e-mail address, make sure to delimit each one with a comma (e.g., jsmith@hotmail.com, jsmith@gmail.com, jsmith@yahoo.com). The blank space after the comma is optional.</p> <hr/> <p>Note: This field is required when the Notify by Email option has been marked.</p> <hr/>
Subject	<p>Determines the subject line of the e-mail notification. When specifying the subject line, you can use keywords that will be replaced with dynamic data when the e-mail notification is sent. For more information on keywords, please refer to the Notification Keywords section above.</p> <hr/> <p>Note: This field is required when the Notify by Email option has been marked.</p> <hr/>

Name	Description
Body	<p>Determines the body of the e-mail notification. When specifying the body, you can use keywords that will be replaced with dynamic data when the e-mail notification is sent. For more information on keywords, please refer to the Notification Keywords section above.</p> <hr/> <p>Note: This field is required when the Notify by Email option has been marked.</p>
Test Notification	<p>Allows you to send a test e-mail notification. The subject line of this e-mail notification will indicate to the recipient that they are receiving a test notification.</p>
Reset	<p>Restores the default subject line and message.</p>

HTTP POST Request Notification

An HTTP POST request can be sent to a server when an alert's condition has been met. The body of this HTTP POST request can contain dynamic data that identifies the network conditions that triggered the alert. Since you can use headers to define the format for the body of the HTTP POST request, you can serve up data in XML, JSON, name/value pairs, etc. One use for this type of notification is to provide the means through which you can feed your API information sent from our network monitoring system.

Before taking advantage of this type of notification, you will need to configure an HTTP server to support HTTP POST requests to a specific URL. After which, you will need to define how your API will use the information sent from our network monitoring system.

Name	Description
Notify by HTTP Post	<p>Determines whether an HTTP POST request will be sent out when an alert's expression and notification conditions are met.</p>
Url	<p>Indicates the URL to which an HTTP POST request will be sent.</p> <hr/> <p>Note: This field is required when the HTTP Post Request option has been marked.</p>

Name	Description
Headers	<p>Determines the headers that will be passed with the HTTP POST request. The format for each request header is <i>Name:Value</i> (e.g., Content-Type: application/xml). Each request header should be placed on a new line.</p> <p>Make sure that the Content-Type header defines the format for the Body option. Your application server will expect data in the specified format.</p> <hr/> <p>Note: This field is required when the HTTP Post Request option has been marked.</p> <hr/>
Body	<p>Determines the body of the HTTP POST request. When specifying the body, you can use keywords that will be replaced with dynamic data when the e-mail notification is sent. For more information on keywords, please refer to the Notification Keywords section above.</p> <hr/> <p>Note: This field is required when the HTTP Post Request option has been marked.</p> <hr/>
Test Notification	<p>Allows you to send a test HTTP POST request. The status code returned by the server will be indicated directly below this option.</p>
Reset	<p>Restores the default content for the Headers and Body option.</p>

Appendix A

Custom Report Fields

This section defines the fields that can be found in reports generated for the Custom Reports module.

Field	Description
2xx	Indicates the total number of requests or data transferred (MB) for the edge CNAME indicated by the Description column that resulted in a 2xx HTTP status code (e.g., 200 OK).
3xx	Indicates the total number of requests or data transferred (MB) for the edge CNAME indicated by the Description column that resulted in a 3xx HTTP status code (e.g., 302 Found or 304 Not Modified).
4xx	Indicates the total number of requests or data transferred (MB) for the edge CNAME indicated by the Description column that resulted in a 4xx HTTP status code (e.g., 400 Bad Request, 403 Forbidden, or 404 Not Found).
5xx	Indicates the total number of requests or data transferred (MB) for the edge CNAME indicated by the Description column that resulted in a 5xx HTTP status code (e.g., 500 Internal Server Error or 502 Bad Gateway).
Cache Hit %	Indicates the percentage of cacheable requests that were served directly from cache to the requester. <hr/> Note: Please refer to the Cache Hit Ratio report for detailed information on how this statistic is calculated. <hr/>
Cache Hits	Indicates the total number of requests or data transferred (MB) for the edge CNAME indicated by the Description column that resulted in a cache hit (i.e., TCP_EXPIRED_HIT, TCP_HIT, or TCP_PARTIAL_HIT). A cache hit occurs when a cached version of the requested content was found.
Data Transferred (MB)	Indicates the total amount of data transferred, in Megabytes, from our edge servers to HTTP clients (i.e., web browsers) for the edge CNAME indicated by the Description column. The amount of data transferred is calculated by adding HTTP response headers with the response body. As a result, the amount of data transferred for each asset will be greater than its actual file size.
Description	Identifies an edge CNAME by its name (e.g., cdn.mydomain.com).

Field	Description
Hits	Indicates the total number of requests to the edge CNAME identified in the Description column.
Misses	Indicates the total number of requests or data transferred (MB) for the edge CNAME indicated by the Description column that resulted in a cache miss (i.e., TCP_CLIENT_REFRESH_MISS, TCP_EXPIRED_MISS, or TCP_MISS). A cache miss occurs when the requested content was not cached on the edge server that honored the request.
No Cache	Indicates the total number of requests or data transferred (MB) for the edge CNAME indicated by the Description column that resulted in a CONFIG_NOCACHE cache status code.
Other	Indicates the total number of requests or data transferred (MB) for the edge CNAME indicated by the Description column that resulted in a HTTP status code that falls outside of the 2xx - 5xx range.
Platform	Indicates the platform that handled the edge CNAME's traffic.
Unassigned	Indicates the total number of requests or data transferred (MB) for the edge CNAME indicated by the Description column for which cache status code or HTTP status code information was not logged.
Uncacheable	Indicates the total number of requests or data transferred (MB) for the edge CNAME indicated by the Description column that resulted in an UNCACHEABLE cache status code.

Advanced Content Analytics Fields

This section defines the fields that can be found in reports generated for the Advanced Content Analytics module.

Reminder: A report may include metrics for "Europe" or the "Asia/Pacific Region." Those items are not meant to provide statistical information on all IP addresses in those regions. Rather, they only apply to requests that originate from IP addresses that are spread out over Europe or Asia/Pacific instead of to a specific city or country.

Field	Definition
% of Complete Downloads	Indicates the percentage of requests for the specified asset that were downloaded to completion. This field is a percentage of the total requests for the specified asset that occurred over the time span covered by the report.
% of Data Transferred	Indicates the percentage of total data transferred from our edge servers to HTTP clients (i.e., web browsers) for requests that match the report's main criteria (e.g., Country, State, File, etc.). This field is a percentage of the total amount of data transferred over the time span covered by the report.
% of Hits	Indicates the percentage of requests that correspond to the report's main criteria (e.g., Country, State, File, etc.). For most reports, this field is a percentage of the total requests that occurred over the time span covered by the report. <hr/> Note: For the By 404 Errors report, this field is the percentage of the total requests that resulted in a 404 Not Found status code. <hr/>
404 hits	Indicates the total number of requests that resulted in a 404 Not Found status code for the asset identified in the File Name column of the By 404 Errors report.
Avg. Playtime	Indicates the average amount of time that users streamed content for requests that match the report's main criteria (e.g., Country, State, File, etc.).
Browser	This field lists the top 250 web browsers used to access CDN content. Each version of a web browser (e.g., Firefox 8.x, Firefox 7.x, and Firefox 3.x, where x indicates a point release) is counted separately by the report.
City	This field lists the top 250 cities from which requests for CDN content originated. A city is identified by its name, state/province, and country abbreviation. If the city or region (i.e., state/province) could not be determined, then it will indicate that they are unknown. If the country is unknown, then two question marks (i.e., ??) will be displayed.

Field	Definition
Complete Downloads	<p>Indicates the number of times that an asset was downloaded completely to a client over the time span covered by the report.</p> <p>In calculating whether a hit is a "complete download," we take the following into consideration:</p> <ul style="list-style-type: none"> • A request must return a 200 OK or a 206 Partial Content status code. • If byte-range requests are detected, then we will ensure that they provide full coverage for the requested asset. • If byte-range requests are detected, then we will sum up the total data transferred for all requests for the same asset that originate from the same client. The amount of data transferred must be equal to or greater than the file size. <hr/> <p>Note: If compression has been enabled on your account, then the recorded file size may be larger than the total bytes transferred. This may lead to inaccurate data.</p> <hr/>
Country	<p>This field lists each country from which CDN content was requested. If the country of origin could not be determined, then it will be reported under the "Unknown" entry.</p>
Data Transferred	<p>Indicates the total amount of data transferred from our edge servers to HTTP clients (i.e., web browsers). The amount of data transferred is calculated by adding HTTP response headers with the response body. As a result, the amount of data transferred for each asset will be greater than its actual file size.</p> <hr/> <p>Note: This sum is calculated solely from requests that correspond to the report's main criteria (e.g., Country, State, File, etc.).</p> <hr/>
Date	<p>Indicates the date for which CDN activity is being reported. The format for dates is YYYY-MM-DD (e.g., 2015-06-01).</p>
Download Attempts	<p>Indicates the total number of unique requests for the asset specified in the File column. This statistic is calculated by summing the following types of requests:</p> <ul style="list-style-type: none"> • All requests for that asset that result in a 200 OK. • All byte-range requests for the start of that asset (e.g., bytes=0-499) that result in a 206 Partial Content. <hr/> <p>Note: This field does not take into account whether the client downloaded the entire asset.</p> <hr/>

Field	Definition
Extension	<p>Indicates the file name extension and Internet media type (e.g., .html [text/html], .htm [text/html], .aspx [text/html], .js [application/javascript], etc.) for the most requested CDN content.</p> <p>Multiple Internet media types can be associated with a single file name extension. As a result, each unique combination of file name extension and Internet media type is considered a different file type.</p>
File	Please see the definition for File Name (below).
File Name	<p>Indicates the CDN URL path to an asset. The CDN URL path always starts with the content access point (i.e., /yyxxx or /yyxxx/<i>CustomerOrigin</i>). The content access point starts directly after the CDN domain (e.g., wac.0001.edgecastcdn.net) in a CDN URL. In the following sample CDN URL, the relative path is indicated in bold font.</p> <p>http://wpc.0001.edgecastcdn.net/800001/videos/video01.flv</p> <hr/> <p>Note: For the purposes of this field, edge CNAME URLs are converted to their equivalent CDN URLs. This allows an accurate tally for all statistics associated with an asset regardless of the CDN or edge CNAME URL used to request it.</p> <hr/>
Full Directory Path	<p>Indicates the CDN URL path to a directory. The CDN URL path always starts with the content access point (i.e., /yyxxx or /yyxxx/<i>CustomerOrigin</i>). The content access point starts directly after the CDN domain (e.g., wac.0001.edgecastcdn.net) in a CDN URL. In the following sample CDN URL, the relative path to a directory is indicated in bold font.</p> <p>http://wpc.0001.edgecastcdn.net/800001/videos/video01.flv</p> <hr/> <p>Note: For the purposes of this field, edge CNAME URLs are converted to their equivalent CDN URLs. This allows an accurate tally for all statistics associated with a directory regardless of the CDN or edge CNAME URL used to request it.</p> <hr/>
Hits	Indicates the total number of requests for CDN content that correspond to the report's main criteria (e.g., Country, State, File, etc.). For the purposes of this field, a hit occurs whenever CDN content is requested, regardless of the status code returned to the client.
Hour	This field lists each hour contained within the By Hour report. Each hour is indicated using the following date/time format: YYYY-MM-DD hh:mm (e.g., 2015-06-01 01:00). Time is reported using 24-hour notation in UTC/GMT time.
Province	This field lists each province from which CDN content was requested. If the origin could not be determined, then it will be reported under the "unknown region" entry.

Field	Definition
Rank	The regions reported in a map-based geography report (e.g., World Map, United States Map, Canada Map, etc.) are listed in order by the number of hits that originated in each region. Regions are listed from the most number of hits to the least number of hits. This field assigns a number to each region based on this order. The region with the most number of hits is assigned the number 1; the remaining regions are numbered sequentially in order.
Referrer	This field lists the top 250 referrers that your clients used to request CDN content. A referrer indicates the URL from which a request was generated. This information is typically stored in an HTTP request header. A dash (-) referrer indicates that the content was hit directly, the referrer was stripped out by the user agent, or the referrer was not passed through the HTTP header.
State	This field lists each state from which CDN content was requested. If the origin could not be determined, then it will be reported under the "unknown region" entry.

Edge Performance Analytics Fields

The fields included in Edge Performance Analytics reports are defined in this section.

Common Edge Performance Analytics Fields

The majority of the Edge Performance Analytics reports have many fields in common. Although these reports share many of the same fields, the data that will be reported for those fields will vary according to the type of report. For example, the Daily Summary report provides statistical information for each day included in the report. This means that the statistics returned for that report will only be for a particular date. Likewise, the Origins report provides statistical information for each customer origin server that served traffic over a specified time period. For this report, field data will be restricted to the traffic served on each customer origin.

Reminder: Both request/response headers and bodies are included when calculating fields that involve the transfer of data.

Field	Description
Hits	Indicates the total number of requests for CDN content that correspond to the criteria specified in the report's first column (e.g., Day, Hour, Protocol, etc.). For the purposes of this field, a hit occurs whenever CDN content is requested, regardless of the status code returned to the client.
Percentage	Indicates the percentage of requests that correspond to the criteria specified in the report's first column (e.g., Day, Hour, Protocol, etc.). This field is a percentage of the total requests that occurred over the time span covered by the report.
Ave Bytes Out	Indicates the average amount of data that was sent from our edge servers to HTTP clients (i.e., web browsers). This average value is calculated solely from the requests that correspond to the criteria specified in the report's first column (e.g., Day, Hour, Protocol, etc.).
Total Bytes Out	Indicates the total amount of data that was sent from our edge servers to your HTTP clients. This sum is calculated solely from the requests that correspond to the criteria specified in the report's first column (e.g., Day, Hour, Protocol, etc.).
Ave Bytes In	Indicates the average amount of data that was received by our edge servers from HTTP clients (i.e., web browsers). This average value is calculated solely from the requests that correspond to the criteria specified in the report's first column (e.g., Day, Hour, Protocol, etc.).
Total Bytes In	Indicates the total amount of data that was received by our edge servers from HTTP clients (i.e., web browsers). This sum is calculated solely from the requests that correspond to the criteria specified in the report's first column (e.g., Day, Hour, Protocol, etc.).
Ave Bytes Remote	Indicates the average amount of data that was sent from CDN and customer origin servers to our edge servers. This average value is calculated solely from the requests that correspond to the criteria specified in the report's first column (e.g., Day, Hour, Protocol, etc.).
Total Bytes Remote	Indicates the total amount of data received from CDN and customer origin servers by our edge servers. This average value is calculated solely from the requests that correspond to the criteria specified in the report's first column (e.g., Day, Hour, Protocol, etc.).
Transfer Rate	Indicates the rate (e.g., Gigabits per second, Megabits per second, Kilobits per second, etc.) at which our edge servers delivered content to HTTP clients (i.e., web browsers). This rate is calculated solely from the requests that correspond to the criteria specified in the report's first column (e.g., Day, Hour, Protocol, etc.).

Field	Description
Duration	Indicates the average amount of time, in milliseconds, that it took our CDN to serve an asset to a client. This average value is calculated solely from the requests that correspond to the criteria specified in the report's first column (e.g., Day, Hour, Protocol, etc.).

Edge Performance Analytics Fields

This section describes all fields used by the reports in the Edge Performance Analytics module.

Note: Definitions for common fields can be found in the **Common Edge Performance Analytics Fields** section above.

Field	Description
403 Hits	Indicates the total number of requests for the asset indicated by the File Name column that resulted in a 403 Forbidden response code.
404 Hits	Indicates the total number of requests for the asset indicated by the File Name column that resulted in a 404 Not Found response code.
4XX Hits	Indicates the total number of requests for the asset indicated by the File Name column that resulted in a 4xx response code.
502 Hits	Indicates the total number of requests for the asset indicated by the File Name column that resulted in a 502 Bad Gateway response code.
503 Hits	Indicates the total number of requests for the asset indicated by the File Name column that resulted in a 503 Service Unavailable response code.
504 Hits	Indicates the total number of requests for the asset indicated by the File Name column that resulted in a 504 Gateway Timeout response code.
5XX Hits	Indicates the total number of requests for the asset indicated by the File Name column that resulted in a 5xx response code.
Ave Bytes In	Please refer to the Common Edge Performance Analytics Fields section above.
Ave Bytes Out	Please refer to the Common Edge Performance Analytics Fields section above.
Ave Bytes Remote	Please refer to the Common Edge Performance Analytics Fields section above.
Cache Status	This field lists the different types of cache statuses returned by the report. A list of the available cache statuses and their definitions is provided in Appendix C: Cache Statuses .
Client	This field lists the IP addresses for the top 100 clients to request content from your CDN account.

Field	Description
Client Request Type	<p>This field lists the types of requests generated by clients for your CDN account. The type of request that will be generated by a client is determined by the header information (e.g., If-Modified-Since or Cache-Control: no-cache) included with the request for the asset.</p> <p>The types of client requests that are tracked are normal, refresh, revalidation, and revalidation + refresh. Each type of client request is defined in the Client Request Types Report section of the Edge Performance Analytics chapter.</p>
CName	<p>This field lists the CNAMEs for which CDN activity took place during the time period covered by the report. This field includes all edge CNAMEs and the CDN domain (e.g., wac.0001.edgecastcdn.net). The CDN domain is included in this list, since it is a CNAME to a server on our network.</p> <hr/> <p>Reminder: Variations in the hostname (i.e., due to case or port) will cause a separate entry to be listed in this field.</p> <hr/>
Compression Type	<p>This field lists the type of compression used to deliver an asset to a client. The available compression types are bzip2, deflate, gzip, not compressible, and uncompressed. Each compression type is defined in the Compression Types Report section of the Edge Performance Analytics chapter.</p>
Day	<p>Indicates the date (YYYY-MM-DD) for which data is being reported.</p>
Duration	<p>Please refer to the Common Edge Performance Analytics Fields section above.</p>
File Name	<p>Indicates the CDN URL path to an asset. The CDN URL path always starts with the content access point (e.g., /000001/). The content access point starts directly after the CDN domain (e.g., wac.0001.edgecastcdn.net).</p> <hr/> <p>Note: For the purposes of this field, edge CNAME URLs are converted to their equivalent CDN URLs. This allows an accurate tally for all statistics associated with an asset regardless of the CDN or edge CNAME URL used to request it.</p> <hr/>
File Type	<p>This field lists the file name extension (e.g., .swf, .png, .mp4, etc.) and the Internet media type associated with the requested content. This field is formatted as follows: <i>.ext [MediaType]</i>.</p>
Geo POP	<p>Indicates the three-letter abbreviation for each POP that handled CDN traffic.</p>
Hits	<p>Please refer to the Common Edge Performance Analytics Fields section above.</p>

Field	Description
Hour	This field lists each hour contained within the Hourly report. Each hour is indicated using the following date/time format: YYYY-MM-DD hh:mm (e.g., 2015-06-01 01:00). Time is reported using 24-hour notation in UTC/GMT time.
HTTP Method	This field lists the HTTP request methods used to request CDN content. The most common HTTP request methods for this report are GET, HEAD, and POST.
HTTP Response Code	This field lists the HTTP response codes (e.g., 200, 403, 404, etc.) that were returned in response to a request for CDN content.
Origin	This field lists the CDN or customer origin servers that hosted CDN content that was requested.
Percentage	Please refer to the Common Edge Performance Analytics Fields section above.
Protocol	This field lists each protocol (i.e., HTTP or HTTPS) used to request CDN content.
Referrer	<p>This field lists the top 100 referrers that your clients used to request CDN content. A referrer indicates the URL from which a request was generated. This information is typically stored in an HTTP request header.</p> <hr/> <p>A dash (-) referrer indicates that the content was hit directly, the referrer was stripped out by the user agent, or the referrer was not passed through the HTTP header.</p> <hr/>
Response Type	Indicates a Token-Based Authentication status code. This status code indicates how a request for an asset was handled by Token-Based Authentication. Each available response type is defined in the Token Auth Summary Report section of the Edge Performance Analytics chapter.
Total Bytes	Indicates the total amount of data transferred for the first instance of each unique asset on that date.
Total Bytes In	Please refer to the Common Edge Performance Analytics Fields section above.
Total Bytes Out	Please refer to the Common Edge Performance Analytics Fields section above.
Total Bytes Remote	Please refer to the Common Edge Performance Analytics Fields section above.
Transfer Rate	Please refer to the Common Edge Performance Analytics Fields section above.

Field	Description
Unique Files	<p>Indicates the total number of unique assets that were requested on a particular date. The URL used to access the asset determines whether an asset is unique. If query-string caching has been enabled, then the query string will also be taken into account when determining whether an asset is unique.</p> <hr/> <p>Note: For the purposes of this field, edge CNAME URLs are converted to their equivalent CDN URLs. This allows an accurate tally for all statistics associated with an asset regardless of the CDN or edge CNAME URL used to request it.</p> <hr/>
URL	<p>Indicates the CDN URL path to an asset. The CDN URL path always starts with the content access point (i.e., /yyxxx or /yyxxx/<i>CustomerOrigin</i>). The content access point starts directly after the CDN domain (e.g., wac.0001.edgecastcdn.net) in a CDN URL. In the following sample CDN URL, the relative path is indicated in bold font.</p> <p>http://wpc.0001.edgecastcdn.net/800001/videos/video01.flv</p> <hr/> <p>Note: For the purposes of this field, edge CNAME URLs are converted to their equivalent CDN URLs. This allows an accurate tally for all statistics associated with an asset regardless of the CDN or edge CNAME URL used to request it.</p> <hr/>
User Agent	<p>Indicates the user agents through which CDN content was requested. Typically, this field will indicate the web browser and version used by the client to request CDN content. Each unique combination of web browser and version number is considered a different user agent.</p>

Appendix B

Monitoring Criteria (Metrics)

A core component of a real-time alert configuration is the metric that will be used to define when an alert notification should be sent out. A metric is a real-time statistic that provides insight into the customer experience. It reveals how much and whether content is being delivered to your customers.

Note: Several metrics report statistical data for a particular cache status. A brief description is provided for each of those metrics below. For a more detailed explanation, please refer to **Appendix C: Cache Statuses**.

Name	Description
Bandwidth Mbps	This metric measures the amount of bandwidth usage for the specified platform in Megabits per second (Mbps).
Cache Status: CONFIG_NOCACHE per second	This metric measures the number of requests per second that result in a CONFIG_NOCACHE status. This status indicates that a customer-specific configuration on our edge servers prevented the asset from being cached.
Cache Status: CONFIG_NOCACHE percentage	This metric measures the percentage of requests per second that result in a CONFIG_NOCACHE status. This status indicates that a customer-specific configuration on our edge servers prevented the asset from being cached.
Cache Status: None per second	This metric measures the number of requests per second that result in NONE status. This status indicates that a cache content freshness check was not performed.
Cache Status: None percentage	This metric measures the number of requests per second that result in NONE status. This status indicates that a cache content freshness check was not performed.
Cache Status: TCP_CLIENT_REFRESH_MISS per second	This metric measures the number of requests per second that result in a TCP_CLIENT_REFRESH_MISS status. This status is reported when an HTTP client (e.g., browser) forces an edge server to retrieve a new version of a stale asset from the origin server.

Name	Description
Cache Status: TCP_CLIENT_REFRESH_MISS percentage	This metric measures the percentage of requests per second that result in a TCP_CLIENT_REFRESH_MISS status. This status is reported when an HTTP client (e.g., browser) forces an edge server to retrieve a new version of a stale asset from the origin server.
Cache Status: TCP_EXPIRED_HIT per second	This metric measures the number of requests per second that result in a TCP_EXPIRED_HIT status. This status is reported when a request that targeted an asset with an expired time to live (TTL), such as when the asset's max-age has expired, was served directly from the POP to the client.
Cache Status: TCP_EXPIRED_HIT percentage	This metric measures the percentage of requests per second that result in a TCP_EXPIRED_HIT status. This status is reported when a request that targeted an asset with an expired time to live (TTL), such as when the asset's max-age has expired, was served directly from the POP to the client.
Cache Status: TCP_EXPIRED_MISS per second	This metric measures the number of requests per second that result in a TCP_EXPIRED_MISS status. This status is reported when a newer version of an expired cached asset is served from the POP to the client.
Cache Status: TCP_EXPIRED_MISS percentage	This metric measures the percentage of requests per second that result in a TCP_EXPIRED_MISS status. This status is reported when a newer version of an expired cached asset is served from the POP to the client.
Cache Status: TCP_HIT per second	This metric measures the number of requests per second that result in a TCP_HIT status. This status is reported when a request is immediately served from the POP to the client.
Cache Status: TCP_HIT percentage	This metric measures the percentage of requests per second that result in a TCP_HIT status. This status is reported when a request is immediately served from the POP to the client.
Cache Status: TCP_MISS per second	This metric measures the number of requests per second that result in a TCP_MISS status. This status indicates that a cached version of the requested asset was not found on the POP closest to the client. These types of requests are forwarded to an origin shield server or an origin server.
Cache Status: TCP_TCP_MISS percentage	This metric measures the percentage of requests per second that result in a TCP_MISS status. This status indicates that a cached version of the requested asset was not found on the POP closest to the client. These types of requests are forwarded to an origin shield server or an origin server.

Name	Description
Cache Status: Total Hits per second	This metric measures the number of successful requests per second for content on the specified platform. For the purposes of this metric, a successful request consists of cache hits and cache misses.
Cache Status: UNCACHEABLE per second	This metric measures the number of requests per second that result in an UNCACHEABLE status. This status is reported when an asset's Cache-Control and Expires headers indicate that it should not be cached on a POP or by the HTTP client.
Cache Status: UNCACHEABLE percentage	This metric measures the percentage of requests per second that result in an UNCACHEABLE status. This status is reported when an asset's Cache-Control and Expires headers indicate that it should not be cached on a POP or by the HTTP client.
Status Code: 2XX per second	This metric measures the number of 2xx status codes (i.e., 200, 201, 202, etc.) that occur per second. This type of status code indicates that the request was successfully delivered to the client.
Status Code: 2XX percentage	This metric measures the percentage of 2xx status codes (i.e., 200, 201, 202, etc.) that occur per second. This type of status code indicates that the request was successfully delivered to the client.
Status Code: 304 per second	This metric measures the number of 304 status codes that occur per second. This status code indicates that the requested asset has not been modified since it was last retrieved by the HTTP client.
Status Code: 304 percentage	This metric measures the percentage of 304 status codes that occur per second. This status code indicates that the requested asset has not been modified since it was last retrieved by the HTTP client.
Status Code: 3XX per second	This metric measures the number of 3xx status codes (i.e., 300, 301, 302, etc.) that occur per second. This type of status code indicates that the request resulted in a redirection.
Status Code: 3XX percentage	This metric measures the percentage of 3xx status codes (i.e., 300, 301, 302, etc.) that occur per second. This type of status code indicates that the request resulted in a redirection.
Status Code: 400 per second	This metric measures the number of 400 status codes that occur per second. This status code indicates that the server refused to process the request due to a client error.
Status Code: 400 percentage	This metric measures the percentage of 400 status codes that occur per second. This status code indicates that the server refused to process the request due to a client error.

Name	Description
Status Code: 401 per second	This metric measures the number of 401 status codes that occur per second. This status code indicates that the request lacked the credentials required for authorization.
Status Code: 401 percentage	This metric measures the percentage of 401 status codes that occur per second. This status code indicates that the request lacked the credentials required for authorization.
Status Code: 403 per second	This metric measures the number of 403 status codes that occur per second. This status code indicates that the request was deemed unauthorized. One possible cause for this status code is when an unauthorized user requests an asset protected by Token-Based Authentication.
Status Code: 403 percentage	This metric measures the percentage of 403 status codes that occur per second. This status code indicates that the request was deemed unauthorized. One possible cause for this status code is when an unauthorized user requests an asset protected by Token-Based Authentication.
Status Code: 404 per second	This metric measures the number of 404 status codes that occur per second. This status code indicates that the requested asset could not be found.
Status Code: 404 percentage	This metric measures the percentage of 404 status codes that occur per second. This status code indicates that the requested asset could not be found.
Status Code: 409 per second	This metric measures the number of 409 status codes that occur per second. This status code indicates that the requested action conflicts with the state of the requested content.
Status Code: 409 percentage	This metric measures the percentage of 409 status codes that occur per second. This status code indicates that the requested action conflicts with the state of the requested content.
Status Code: 429 per second	This metric measures the number of 429 status codes that occur per second. This status code indicates that the client submitted too many requests within a short time period.
Status Code: 429 percentage	This metric measures the percentage of 429 status codes that occur per second. This status code indicates that the client submitted too many requests within a short time period.
Status Code: 4XX per second	This metric measures the number of 4xx status codes (i.e., 400, 401, 402, 405, etc.) that occur per second. This status code indicates that the requested asset was not delivered to the client.

Name	Description
Status Code: 4XX percentage	This metric measures the percentage of 4xx status codes (i.e., 400, 401, 402, 405, etc.) that occur per second. This status code indicates that the requested asset was not delivered to the client.
Status Code: 5XX per second	This metric measures the number of 5xx status codes (i.e., 500, 501, 502, etc.) that occur per second. This status code indicates that the server was unable to honor the request.
Status Code: 5XX percentage	This metric measures the percentage of 5xx status codes (i.e., 500, 501, 502, etc.) that occur per second. This status code indicates that the server was unable to honor the request.
Status Code: Other per second	This metric measures the total occurrences for all status codes that are not covered by another metric.
Status Code: Other percentage	This metric measures the percentage of requests that returned a status code that was not covered by another metric.
Status Code: Total Hits per second	This metric measures the number of requests per second. You can use this option as a baseline indicator to see the percentage of total hits that a particular status code comprises.
Total Connections	This metric measures the number of new connections per second.

Appendix C

Cache Statuses

Each cache status that is reported for CDN activity is defined below.

Cache Status	Description
CONFIG_NOCACHE	This status indicates that a customer-specific configuration on our edge servers prevented the asset from being cached. For example, an HTTP Rules Engine rule can prevent an asset from being cached by enabling the Bypass Cache feature for qualifying requests.
NONE	This status indicates that a cache content freshness check was not performed. This check is skipped when Token-Based Authentication denies a request or when an HTTP request method is used that bypasses cache (e.g., PUT, DELETE, etc).
TCP_CLIENT_REFRESH_MISS	<p>This status is reported when an HTTP client (e.g., browser) forces an edge server to retrieve a new version of a stale asset from the origin server.</p> <p>By default, our servers prevent an HTTP client from forcing our edge servers to retrieve a new version of the asset from the origin server. However, this behavior can be overridden through the use of the HTTP Rules Engine feature called "Honor No-Cache Request."</p>
TCP_EXPIRED_HIT	This status is reported when a request that targeted an asset with an expired time to live (TTL), such as when the asset's max-age has expired, was served directly from the POP to the client. An expired request typically results in a revalidation request to the origin server. In order for a TCP_EXPIRED_HIT to occur, the origin server must indicate that a newer version of the asset does not exist. This type of situation will typically update that asset's Cache-Control and Expires headers.
TCP_EXPIRED_MISS	This status is reported when a newer version of an expired cached asset is served from the POP to the client. This occurs when the TTL for a cached asset has expired (e.g., expired max-age) and the origin server returns a newer version of that asset. This new version of the asset will be served to the client instead of the cached version. Additionally, it will be cached on the edge server and the client.

Cache Status	Description
TCP_HIT	This status is reported when a request is immediately served from the POP to the client. An asset is immediately served from a POP when it is cached on the POP closest to the client and it has a valid TTL. TTL is determined by the Cache-Control: s-maxage, Cache-Control: max-age, and Expires headers.
TCP_MISS	This status indicates that a cached version of the requested asset was not found on the POP closest to the client. The asset will be requested from either an origin server or an origin shield server. If the origin server or the origin shield server returns an asset, it will be served to the client and cached on both the client and the edge server. Otherwise, a non-200 status code (e.g., 403 Forbidden, 404 Not Found, etc.) will be returned.
TCP_PARTIAL_HIT	<p>This status is reported when a request results in a hit for a partially cached asset. The requested asset is immediately served from the POP to the client.</p> <hr/> <p>Note: The Partial Cache Sharing feature (HTTP Rules Engine) enables the capability to generate partially cached content.</p> <hr/>
UNCACHEABLE	This status is reported when an asset's Cache-Control and Expires headers indicate that it should not be cached on a POP or by the HTTP client. These types of requests are served from the origin server.

Glossary

A

Asset

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

C

Cache

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

CDN

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

CDN Domain

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

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

CDN Origin

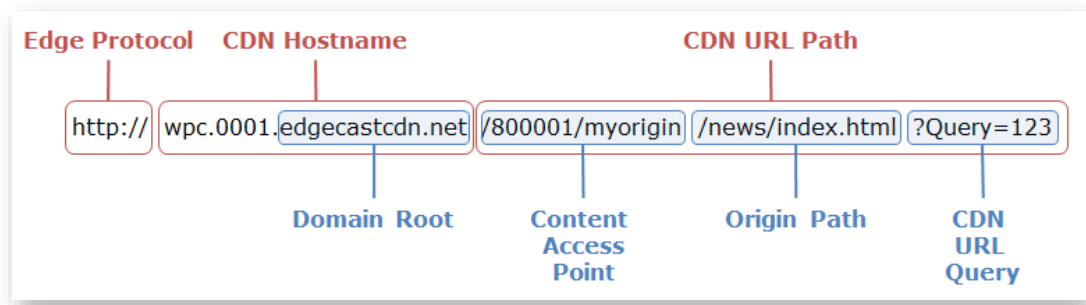
This term refers to a storage server on our CDN. Our CDN origin servers are in close proximity to our POPs, in order to provide optimal conditions for transferring data from a CDN origin server to your end-users via our POPs.

CDN Origin Identifier

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

CDN URL

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



CDN URL Path

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

- `/00xxxx/sales/brochures/widgets.pdf`
- `/00xxxx/tech/whitepapers/sprockets.pdf`
- `/00xxxx/marketing/videos/presentation01.flv`

CDN/Edge CNAME URL Query

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

CNAME

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

Content Access Point

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

Customer Origin

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

Customer Origin Identifier

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

D

Domain Root

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

Dynamic Streaming

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

E

Edge CNAME

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

Edge CNAME URL

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

In the following example, the domain assigned to the edge CNAME is "cdn.mydomain.com."

Edge CNAME URL	Points To
http://cdn.mydomain.com/marketing/videos/presentation01.flv	http://wpc.xxx.edgecastcdn.net/00xxx/marketing/videos/presentation01.flv

Edge CNAME URL Path

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

Edge CNAME URL Path	Actual CDN URL
/videos	http://wpc.xxx.edgecastcdn.net/00xxx/marketing/videos/
/videos/Show01.flv	http://wpc.xxx.edgecastcdn.net/00xxx/marketing/videos/Show01.flv

Edge Protocol

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

Edge Server

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

Encryption Key

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

G

Global Key

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

H

HTTP Large

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

HTTP Progressive Download

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

HTTP Small

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

I

Identifier

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

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

Ingest

This term refers to the process of capturing and transforming video into a stream.

Ingest Server

This term refers to the type of server that is dedicated to the process of capturing and transforming media into a stream. This type of server will then broadcast that stream throughout our CDN.

L

Live Authentication Key

This type of key authenticates a stream before it is ingested by our publishing server. There are two types of Live Authentication keys, which are global and stream keys. A live authentication key must be specified when setting an encoder's stream setting. The proper notation is provided below.

```
StreamName?LiveAuthenticationKey
```

Live Ingestion Point

This term refers to the location on a server where our CDN can access encoded media. Our streaming services can ingest live streams via a publishing point.

Load

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

M

Media Control Center (MCC)

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

<https://my.edgecast.com>

O

Origin Path

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

Origin Server

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

Origin Shield

This feature provides a layer of protection for your customer origin server by creating an intermediate caching layer between it and our edge servers. This caching layer resides on one or more of our point-of-presence (POPs). Requests that have not been previously cached on a POP will be channeled through the closest origin shield server. The origin shield server will then either serve a cached version of the requested content or retrieve it from your customer origin server. This feature reduces the amount of bandwidth used on your customer origin server, since most requests will be handled by the origin shield server.

P

Player URL

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

Point-of-Presence (POP)

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

Pre-Cached

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

Publishing Point

This term refers to the location on the publishing server to which your encoder will broadcast encoded media.

Publishing Server

This term refers to a CDN server that will redistribute encoded media as a streamer that will be broadcast to your end-users via our POPs.

Purge

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

Push Stream

This type of stream requires your encoder to send, or push, encoded video to a CDN server. From there, our server will create a stream and deliver it to clients that request it.

Q

Query String

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

R

Request

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

Response

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

S

Stream

A stream consists of the delivery of audio/video content in a format that allows your clients to play it back through a multimedia player.

Stream Key

This type of Live Authentication key can only authenticate a stream when it is published to the path associated with it.

T

Time to Live (TTL)

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

Token

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

Token-Based Authentication

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