



Media Streamer Release 3.14 User Guide

First Published: 2017-10-01

Last Modified: 2019-05-09

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2017–2019 Cisco Systems, Inc. All rights reserved.



CONTENTS

Preface **xiii**

Audience **xiii**

Conventions **xiii**

Communications, Services, and Additional Information **xiv**

CHAPTER 1

CDN and Cisco Media Streamer Overview **1**

What is a CDN? **1**

Cisco Media Streamer Overview **1**

Key Media Streamer Terminology **3**

Key Functions of Media Streamer **3**

CHAPTER 2

Cisco OMD Director Overview **7**

Logging into Cisco OMD Director **7**

 OMD Director HA **8**

 Log in to OMD Director **8**

 Logging out of OMD Director **10**

OMD Director Portal **10**

OMD Director User Interface Elements **10**

OMD Director Navigation Panel: KPI Metrics Overview **12**

 KPI Metrics > CDN Edge **12**

 KPI Metrics > CDN Mid **14**

OMD Director Navigation Panel: Monitor Overview **16**

 Monitor > Server Metrics **16**

 Monitor > Alarms **16**

OMD Director Navigation Panel: Provisioning Overview **17**

OMD Director Navigation Panel: Insights Overview **18**

OMD Director Navigation Panel: Administration	19
OMD Director Navigation Panel: Support	20
OMD Director Navigation Panel: Broadcasting Groups	22

CHAPTER 3

Initial CDN Provisioning 23

Understanding CDN Client Blocking Options	23
Edge Geo Blocking	26
OMD Director CDN Deployment Steps	27
Using the CDN Wizard	27
Step 1: Create CDN	28
Step 2: Prepare Servers	29
Step 3: Assign Servers	31
Step 4: Create Delivery Services	35
Delivery Service Advanced Settings	38
Step 5: Review	50
Step 6: Accept	51
Validating the CDN Configuration	53
View Delivery Services	54
View Mid Cache Groups	56
View Edge Cache Groups	58
View Servers	60
Print or Export CDN Overview Page	62

CHAPTER 4

Manage CDN Servers 63

Add an Origin Server to the CDN	63
Add a New Registered Server to the CDN	65
Edit Traffic Monitor or Traffic Router Servers	66
Edit an Existing Edge Cache Server	67
Manage Secondary Streaming IPs	68
Edit an Existing Mid Cache Server	70
Edit an Existing Origin Server	71
Delete a Server from CDN Overview Window	72
Delete an Origin Server from Edit CDN Window	73

CHAPTER 5**Manage Cache Groups 75**

- Add a Edge Cache Group 76
- Backup Edge Cache Groups 78
 - Process When Geo Limit is Set to CZF Only 78
 - Process When Geo Limit is Set to CZF + Country Codes or None 79
- Configure Backup Edge Cache Groups 80
- CZF File 82
- Add a Mid Cache Group 84
- Edit an Existing Cache Group 86
- Delete a Mid Cache Group 88
- Delete an Edge Cache Group 89

CHAPTER 6**Manage Delivery Services 91**

- Add a New DNS or HTTP Delivery Service 92
 - Delivery Service Advanced Settings 100
 - Advanced Settings: General 101
 - Advanced Settings: Content Preposition 106
 - Advanced Settings: Header Rewrite Configurations 108
 - Advanced Settings: URL Signing 110
 - Advanced Settings: Traffic Router CORS Settings for HTTP(s) Content Routing 112
 - Advanced Settings: Geo Limit 113
- Clone a Delivery Service 116
- Steering Delivery Service 116
 - Target Delivery Service Selection 117
 - Create a Steering Delivery Service Overview 117
 - Create a Steering Delivery Service 118
 - Assign Target Delivery Service to a Steering Delivery Service 124
- Edit an Existing DNS or HTTP Delivery Service 124
- Edit an Existing Steering Delivery Service 125
- Delete a Delivery Service 126
- Advanced Delivery Service Features 126
 - URL Signing 127
 - Configure Delivery Service Profile for URL Signing 128

Session Tracking Security Enablement	132
Multi Site Origin	133
Multi Site Origin Configuration Overview	134
Configure Multi Site Origin	135
View Origin Server Cache Groups	135
Create Origin Server Cache Groups	136
View Origin Server Profiles	137
Create Origin Server Profile	139
Add Origin Servers to the CDN for MSO	141
Assign Origin Cache Groups to Mid Cache Groups	141
Configure Delivery Service Profile for MSO	142
Configure the Delivery Service for MSO	146
Regex Remap Settings	147
Edge Geo Blocking	150
Configure Edge Geo Blocking	151
CDN Routing Name	154

CHAPTER 7

Manage Client Routing	155
Client Routing Overview	155
Coverage Zone File	156
Managing the CZF File	158
Upload a CZF File	159
Delete a CZF File	159
Proximity Routing	159
Geolocation Based Routing	161

CHAPTER 8

Manage Profiles	163
Profiles Overview	163
Add a Profile	165
Assign a CZF to a Traffic Router Profile	166
Configure Proximity Routing	166
Configure the Proximity Server	168
Anonymous Blocking	169
Configure Anonymous Blocking	169

Upload an Anonymous IP Database	169
Enable and Configure Anonymous Blocking in the Profile	170
Enable Anonymous Blocking on a Delivery Service	173
ASN Blocking	173
Configure ASN Blocking	174
Upload an ISP Database	174
Assign an ISP Database to the Traffic Router Profile	174
Enable ASN Blocking on a Delivery Service	175
NGB Whitelist	176
Configure an NGB Whitelist	176

CHAPTER 9	Manage Content Invalidation Policies	177
	Add a Content Invalidation Policy	177
	Delete a Content Invalidation Policy	179

CHAPTER 10	Device Groups	181
	Device Groups Overview	181
	Create a Device Group	182
	Edit a Device Group	184
	Delete a Device Group	185

CHAPTER 11	KPI Metrics	187
	Getting Started	187
	KPI Metrics > CDN Edge	188
	CDN Topology	189
	Delivery Bandwidth	190
	Cache Efficiency	190
	Cache Server Response Codes	191
	Cache Server Request Rate	192
	Concurrent Sessions	192
	Unique Clients	193
	System Overview	193
	Filtering the CDN Edge Graphs	194
	KPI Metrics > CDN Mid	195

CDN Topology	195
Cache Fill Bandwidth	197
Mid Cache Efficiency	197
Content Ingest Requests	198
Content Ingest Response Codes	198
Filtering the CDN Mid Graphs	199

CHAPTER 12

CDN Monitoring 201

Server Metrics	201
Change Time Range and Interval	203
Shift Time and Zoom Time	204
Mouse Over	204
Alarms	205
View Active Alarms	205
Acknowledge an Alarm	206
Search for an Alarm	207
Filter an Alarm	208
Alarms History	208
Filter Alarms History	209
Configure Alarm Rules	209
Create New Alarm Rule	210
Edit an Alarm Rule	212
Deactivate an Alarm Rule	213
Mute Alarms	213

CHAPTER 13

OMD Insights 215

Insights Page Overview	215
Insights Overview Tab	216
Filtering Content	217
Network Tab	218
Caching Tab	218
Content Tab	219
Viewers Tab	220
Content Origin Tab	220

Content Routing Tab	221
Insights Trends Tab	222
Filtering Content	222
Insights Reports Tab	223
Insights Analytics Tab	224
Insights Monitor Tab	226
Filtering Content	227
Insights Alerts Tab	227
Add an Alert	228
Insights Custom Dashboards Tab	229
Add a Custom Dashboard	229
Delete a Custom Dashboard	232
Insights Search Tab	233

CHAPTER 14

OMD Administration	235
Users/Organizations	236
Users	236
Add a New User	236
Edit an Existing User	239
Delete a User	240
Reset a User Passphrase	240
Organizations	241
View Organizations	242
Add a New Organization	243
Delete an Organization and its Users	244
Backup Management	244
Schedule Backups	245
Change Schedule	246
Enable/Disable a Scheduled Backup	247
Manually Back Up and Restore a Database	247
Backup Settings	248
Add Remote Backup Destinations	248
Edit or Delete a Remote Backup Destination	249
Add a New Database Backup	250

Edit a Database Configuration	251
User Profile	251
Change User Settings	252
Notification Settings	253
Activity Log	255

APPENDIX A	Example CZF File	257
-------------------	-------------------------	------------

APPENDIX B	Example Ingest Manifest File	261
-------------------	-------------------------------------	------------

APPENDIX C	NGB Whitelist File	265
-------------------	---------------------------	------------

APPENDIX D	Header Rewrite Rules Syntax	267
	Conditions	268
	Operators	269
	Header Rewrite Rule Example	269
	Configure Header Rewrite Rules	269

APPENDIX E	Configuring Enhanced DNS Request Routing	271
	Configure Enhanced DNS Request Routing	271

APPENDIX F	Recreate OMD Insights Summary Index Buckets	275
-------------------	--	------------

APPENDIX G	tacreport tool	277
-------------------	-----------------------	------------

APPENDIX H	Access and Transaction Log Details	279
	Traffic Router Log Information	279
	Media Streamer Cache Server Transaction Log Information	285

APPENDIX I	OMD Director Alarms and Remediation	289
-------------------	--	------------

APPENDIX J	OMD Monitor Alarms and Remediation	309
	OMD Monitor Client Checks	309

APPENDIX K	Changing Mongoddb Username and Password for OMD Director	337
-------------------	---	------------

APPENDIX L**Manage Content Invalidation using the OMD Director REST API 339****Authentication 339****Generate Token 340****Content Invalidation 341****List Delivery Services 341****List Content Invalidation Jobs 343****List Specific Content Invalidation Job 344****Create a Content Invalidation Job 345****Delete a Content Invalidation Job 346**

Preface

The *Cisco Open Media Distribution User Guide* provides information on how to use OMD Director to provision, manage, and monitor the Cisco Open Media Distribution (OMD) environment.

This preface describes who should read the *Cisco Open Media Distribution User Guide*, how it is organized, and its document conventions. It contains the following sections:

Audience

This guide is for the networking professionals that manage the Cisco Open Media Distribution solution.

Conventions

This document uses the following conventions.

Convention	Indication
bold font	Commands and keywords and user-entered text appear in bold font.
<i>italic font</i>	Document titles, new or emphasized terms, and arguments for which you supply values are in <i>italic</i> font.
[]	Elements in square brackets are optional.
{x y z }	Required alternative keywords are grouped in braces and separated by vertical bars.
[x y z]	Optional alternative keywords are grouped in brackets and separated by vertical bars.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.
courier font	Terminal sessions and information the system displays appear in courier font.
< >	Nonprinting characters such as passwords are in angle brackets.
[]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

**Note**

Means reader take note. Notes contain helpful suggestions or references to material not covered in the manual.

**Caution**

Means reader be careful. In this situation, you might perform an action that could result in equipment damage or loss of data.

**Warning****IMPORTANT SAFETY INSTRUCTIONS**

Means danger. You are in a situation that could cause bodily injury. Before you work on any equipment, be aware of the hazards involved with electrical circuitry and be familiar with standard practices for preventing accidents. Use the statement number provided at the end of each warning to locate its translation in the translated safety warnings that accompanied this device.

SAVE THESE INSTRUCTIONS

Communications, Services, and Additional Information

- To receive timely, relevant information from Cisco, sign up at [Cisco Profile Manager](#).
- To get the business impact you're looking for with the technologies that matter, visit [Cisco Services](#).
- To submit a service request, visit [Cisco Support](#).
- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit [Cisco Marketplace](#).
- To obtain general networking, training, and certification titles, visit [Cisco Press](#).
- To find warranty information for a specific product or product family, access [Cisco Warranty Finder](#).

Cisco Bug Search Tool

[Cisco Bug Search Tool](#) (BST) is a web-based tool that acts as a gateway to the Cisco bug tracking system that maintains a comprehensive list of defects and vulnerabilities in Cisco products and software. BST provides you with detailed defect information about your products and software.



CHAPTER 1

CDN and Cisco Media Streamer Overview

This chapter provides an overview of what a content delivery network (CDN) is and describes the Cisco Media Streamer solution, which is part of the Cisco Open Media Distribution (OMD) Suite. This chapter includes the following information:

- [What is a CDN?, on page 1](#)
- [Cisco Media Streamer Overview, on page 1](#)
- [Key Media Streamer Terminology, on page 3](#)
- [Key Functions of Media Streamer, on page 3](#)

What is a CDN?

A CDN is a system of distributed servers that deliver content to a user based on the geographic locations of the user, the origin of the requested content, and a content delivery service. The goal of a CDN is to provide efficient delivery of streaming content over HTTP/HTTPS to optimize content delivery and to serve content to end-users with high availability and high performance, while minimizing the traffic on the network.

At its core, a CDN performs two essential functions:

- It caches content at the edge of the network, closer to end users, to reduce the IP video traffic that needs to traverse the core network, delivering a higher quality experience to viewers.
- It positions multiservice, multiprotocol content streaming capabilities at the network edge, allowing the operator to adapt video content for virtually any IP video device close to the user that is consuming it.

By 2020 there will be 11 billion connected video devices, and 82 percent of IP traffic will be video (Cisco Visual Networking Index). To meet that demand, today's CDN infrastructures must evolve to scale cost-effectively, accelerate feature velocity, and deliver simplified and open management tools. By adopting cloud architecture and agile software development methodology and using best-in-class open-source software, Cisco Media Streamer provides an open and flexible CDN platform that delivers the multiscreen Internet video quality that consumers expect and service providers can deploy.

Cisco Media Streamer Overview

The Cisco Media Streamer content delivery platform is designed to deliver immersive multiscreen video experiences to managed and unmanaged devices across telco, cable, and mobile access networks. Media Streamer scales cost-effectively to distribute terabits per second (Tbps) of live, on-demand, and time-shifted

video. It enables service providers to compete with over-the-top (OTT) video offerings and generate revenue from wholesale CDN services within their infrastructure. Media Streamer is the foundational IP delivery platform for Cisco's Infinite Video Platform, which provides comprehensive consumer video experiences.

Media Streamer includes all the core elements of management, request routing, load balancing, caching, and analytics to deliver HTTP and HTTPS content at scale and to easily integrate into your network and middleware. Media Streamer builds on Cisco's more than 10 years of CDN expertise.

Media Streamer caches and delivers web content, software, and streaming media with support for media players using the following HTTP steaming protocols: Apple HTTP Live Streaming (HLS), Microsoft HTTP Smooth Streaming (HSS), Adobe HTTP Dynamic Streaming (HDS), and MPEG Dynamic Adaptive Streaming over HTTP (MPEG-DASH). Media Streamer supports video on demand (VoD), live video, time-shifted TV (TSTV), progressive download, secure download, and small object caching from a common high-performance HTTP cache. Media Streamer performs sophisticated algorithms for cache selection based on client location, cache availability, cache load, and content requested.

The primary components of a Cisco Media Streamer solution are:

- **OMD Director:** Cisco OMD Director is a cloud-based CDN management system that provides integrated provisioning, monitoring, analytics, alerting, and role-based management. Cisco OMD Director is implemented to be virtualized and further optimized with microservices in containers. OMD Director is the primary user interface of Cisco Media Streamer.
- **OMD Director Portal:** An OMD Director portal is optional and does not replace the primary OMD Director Master Controller and Worker nodes. It is intended for customers with downstream Resellers or Content Providers who require access to CDN analytics. An OMD Director portal instance is a separate setup from the primary OMD Director installation, with its own user database. It provides a Director GUI that is limited to working with the CDN analytics dashboards provided by Insights. You cannot manage or view the CDN configuration from an OMD Director portal instance.
- **Media Streamer Core Components:**
 - **OMD Core Traffic Server:** A Traffic Server is an HTTP/S proxy cache that is deployed as either an "edge" or "mid-tier" cache server to form a two-tiered CDN hierarchy. Traffic Servers are fast, scalable, extensible using plug-ins, and HTTP/1.1 compliant.
 - **OMD Core Traffic Router:** A Traffic Router is a CDN load balancer that redirects HTTP(s) client requests to an edge cache to ensure that the end user is connected to the optimal cache. The Traffic Router determines which cache to redirect a request to based on client proximity, cache load, and content affinity. The Traffic Router is authoritative for the CDN domain and it implements DNS and HTTP routing.
 - **OMD Core Traffic Ops:** Traffic Ops is an open source management system used to configure advanced settings of the Media Streamer CDN.
- **OMD Insights:** OMD Insights is an application that provides CDN application level insights including operations, content popularity, user consumption, and quality of service (QoS). OMD Insights is based on Splunk and aggregates the log data from all of the Traffic Servers and Traffic Routers.
- **OMD Monitor:** OMD Monitor is an application that provides in-depth server monitoring, threshold crossing, and alarming based on CPU utilization, port utilization, temperature, disk I/O, and other detailed server metrics.

Key Media Streamer Terminology

This section describes some of the key terms you need to be familiar with when provisioning and managing an Media Streamer CDN.

- **Delivery service:** A software structure in Media Streamer that maps an Origin Server to Traffic Servers using an FQDN. It defines a URL that is used to represent an Origin Server and which cache groups can serve content from that server. The Delivery Service also contains configuration parameters that dictate how content is ingested, distributed, and delivered to client devices.
- **Cache group:** A cache group is a logical grouping of caches (Traffic Servers) used to provide high availability. A cache group has one single set of geographical coordinates even if the caches that make up the cache group are in different physical locations. To provide site-level redundancy, caches in a cache group should be in separate physical locations. Cache groups are defined to contain either edge caches or mid-tier caches. A cache group serves a particular part of the network as defined in the coverage zone file.
- **Cache:** Media Streamer uses a two-tiered CDN Traffic Server (cache) hierarchy:
 - **Edge caches:** Provide edge caching, content streaming, and download to subscriber IP devices. Traffic routers redirect client requests to edge caches based on geolocation, server availability, server load, and server cache content to provide efficient system-wide load balancing. Edge caches are organized into cache groups. Each edge cache group is configured with a single mid-tier parent cache group and optionally a secondary mid-tier parent cache group for failover.
 - **Mid caches:** Provide content ingest and storage functionality. When an edge cache does not contain the content requested by the client, the edge cache will proxy the request to a mid-cache server, based on the parent cache group assigned to the edge cache. If the mid-cache server does not contain the content, it is responsible for fetching the content from the Origin Server. Mid-tier caches are also organized into cache groups. Mid-tier cache groups may serve (be a parent to) multiple edge cache groups.

Key Functions of Media Streamer

The following functions are the key functions that Media Streamer provides:

- **Ingest and Distribution:**
 - **Dynamic:** When an edge cache receives a client request, the cache server checks its local cache for the content. If the edge cache does not contain the content, it requests the content from a mid-tier cache. The mid-tier cache checks its local cache for the content. If the mid-tier cache does not contain the content, it will request the content from the Origin Server.
 - **Preposition:** Typically content is not stored on the cache until a client requests the content. However, in some situations you may want to put content on the cache before it is requested. That is what prepositioning enables you to do. When you configure prepositioning, you define what content to put on the selected caches ahead of time by using a Ingest Manifest file.
- **Delivery:** The Traffic Router handles client requests for content and determines which client requests to allow. The Traffic Router also determines the best Cache Group to deliver the content based on client proximity, cache load, and content affinity.

- The Traffic Router supports the following methods to determine which client requests to allow:
 - Coverage Zone File (CZF)
 - National geoblocking (NGB), using a geolocation database, and an NGB whitelist
 - MaxMind Anonymous IP database, which can be used to block:
 - Anonymous VPNs
 - Hosting Providers
 - Public Proxies
 - Tor Exit nodes



Note A license is required to use the MaxMind Anonymous IP database. For more information on obtaining this license, please contact your Cisco Account team.

- ASN blocking, which uses a MaxMind ISP database to block client requests based on the Autonomous System Number (AS Number) to which the client IP address belongs.



Note A license is required to use the MaxMind ISP database. For more information on obtaining this license, please contact your Cisco Account team.

Which methods are used depends on the options that you configure for the Media Streamer deployment. Please refer to [Understanding CDN Client Blocking Options](#) for more information on these methods and how they work together.

- The Traffic Router supports the following Client Routing methods to determine which cache group is the best to deliver the content:
 - **Coverage Zone File (CZF):** The CZF is a static JSON file that maps IP address ranges to cache groups. The Traffic Router checks the CZF for an IP address range that matches the requesting address to determine the best Cache Group to deliver the content.
 - **Proximity routing:** Proximity routing uses network proximity maps that leverage routing information databases to help determine the best Cache Group to deliver the content.



Note Currently Proximity routing requires the use of VDS-IS Proximity Engines (PxE) and is intended for Media Streamer customers that are migrating from VDS-IS.

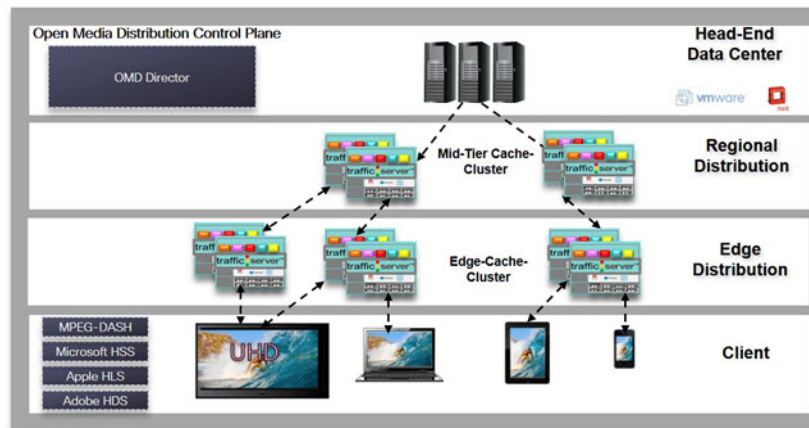
- **Geolocation based routing:** Geolocation based routing uses a geolocation database to determine the best Cache Group to deliver the content.

After the Traffic Router determines the Cache Group to use, it will determine an Edge Cache in that group to use based on cache availability, cache load, and content affinity.

Refer to [Client Routing Overview](#) for more details on each method and how these Client Routing methods interoperate.

- **Management:** OMD Director is the primary Media Streamer management interface. It is a centralized system management application that interacts directly with all of the key components of Media Streamer to enable remote monitoring, management, and diagnostics of the components of the Media Streamer solution from a single tool.

The following diagram shows a typical Cisco Media Streamer deployment.





CHAPTER 2

Cisco OMD Director Overview

Cisco OMD Director is the primary user interface for Cisco Media Broadcaster and Cisco Media Streamer. It provides a unified interface to seamlessly manage the different components of the Cisco OMD deployment. OMD Director enables the administrator to rapidly provision new CDNs and caches, and integrates monitoring, alerts, and real-time analytics into a single management tool.

This chapter walks through the features of Cisco OMD Director to provide a high-level overview of the product and describes the menu options that will be used to manage Media Streamer. Each feature will be explored in more detail in later chapters.



Note

Some of the OMD Director menus apply only to Media Broadcaster functions. This guide will not cover those menus. For more information on managing a Media Broadcaster deployment, see the *Cisco Media Broadcaster User Guide*.

This chapter includes the following topics:

- [Logging into Cisco OMD Director, on page 7](#)
- [OMD Director Portal, on page 10](#)
- [OMD Director User Interface Elements, on page 10](#)
- [OMD Director Navigation Panel: KPI Metrics Overview, on page 12](#)
- [OMD Director Navigation Panel: Monitor Overview, on page 16](#)
- [OMD Director Navigation Panel: Provisioning Overview, on page 17](#)
- [OMD Director Navigation Panel: Insights Overview, on page 18](#)
- [OMD Director Navigation Panel: Administration, on page 19](#)
- [OMD Director Navigation Panel: Support, on page 20](#)
- [OMD Director Navigation Panel: Broadcasting Groups, on page 22](#)

Logging into Cisco OMD Director

The following sections describe how to access the OMD Director GUI. To access the OMD Director GUI, you must be using one of the following supported browsers:

- Firefox version 59 or later
- Chrome version 66 or later

OMD Director HA

OMD Director supports an Active/Standby High Availability (HA) configuration and supports inter-site redundancy. OMD Director HA provides the following:

- Prevents the loss of any configuration data due to a hardware, software, or network failure
- Minimizes the loss of monitoring data such as time series metrics, alarms, and notifications due to a hardware, software, or network failure
- Makes sure alarm emails continue to be sent
- Continues to support streaming capabilities

In an OMD Director HA configuration, your OMD Director instance can be running in either primary or backup mode. If either the OMD Director primary or backup instance loses connection or fails, the remaining OMD Director instances will transition into a detached state.

The following describes the actions you can perform depending on the mode of the OMD Director instance that you are logged into:

- **Primary:** You can only make configuration changes from the primary OMD Director instance. The primary instance is the authoritative instance for all data. Changes to all configurations made from the primary OMD Director instance are replicated to the backup instance.
- **Backup:** While logged into the OMD Director instance that is running in backup mode, you can only view the configuration, you cannot make any changes to the configuration. Any buttons or links that would allow configuration changes will be disabled on the backup OMD Director instance. If you are logged into the backup OMD Director instance, the title bar will show “This OMD Director is currently running in backup mode”.
- **Detached:** While an OMD Director instance is running in a detached state, you can view the configuration and you can acknowledge, clear, and comment on alarms. The buttons or links for any actions that are not allowed will be disabled. If you are logged into an OMD Director instance that is running in the detached state, the title bar will show “This OMD Director is currently running in detached mode”.

If the primary OMD Director instance fails, you must manually configure the backup OMD Director instance as the new primary OMD Director instance. For information on this procedure, see the *Cisco Media Streamer and Cisco Media Broadcaster Installation and Upgrade Guide* for more information on this procedure.

Log in to OMD Director

Perform the following steps to log in to OMD Director:



Note

If you are running OMD Director in an HA configuration, make sure you are logging into the OMD Director Primary instance if you want to make changes to the CDN configuration.

Procedure

Step 1 From a web browser, use HTTPS and enter the hostname or IP address of the OMD Director and port number 8099. For example:

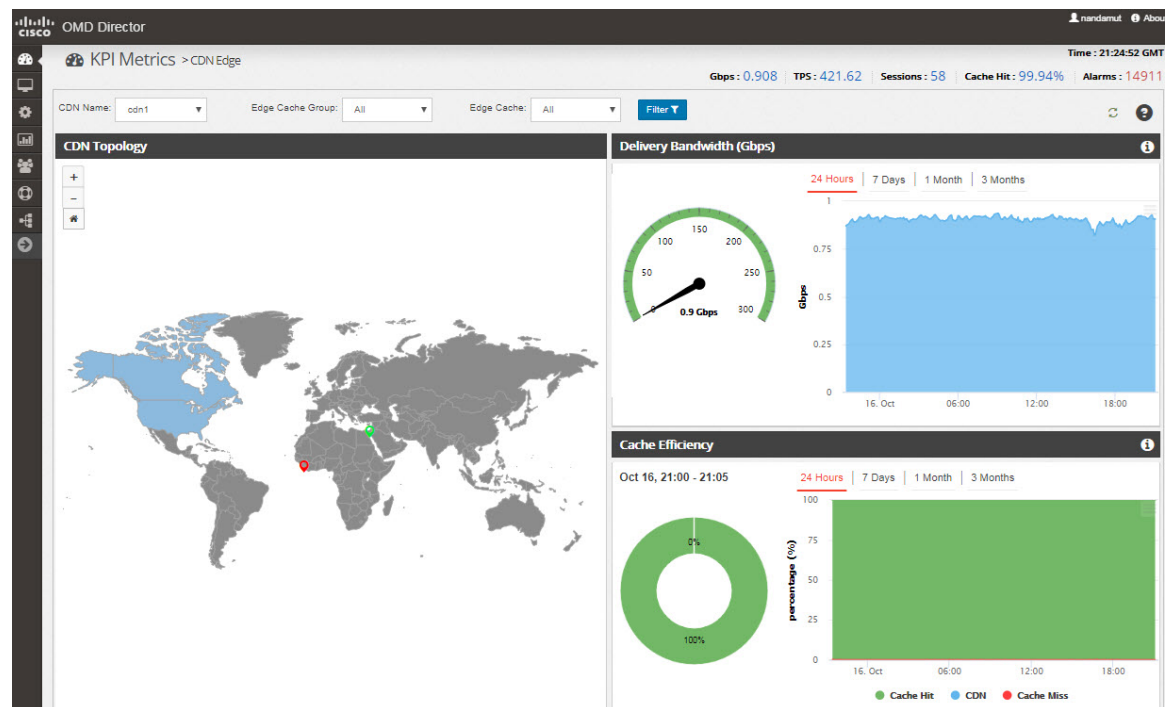
https://<omd_director_worker_hostname_or_ip_address>:8099

The OMD Director login page is displayed.

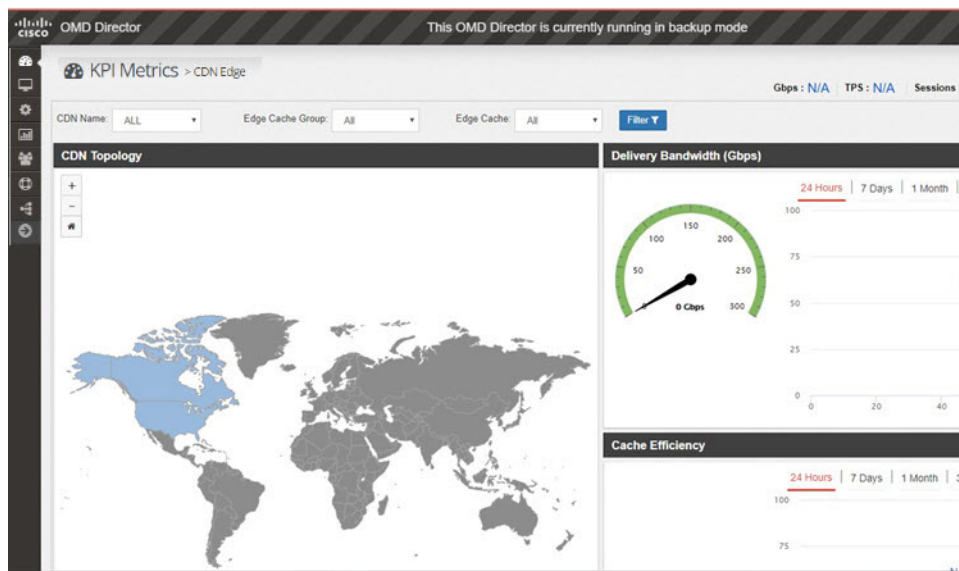
Step 2 Log in with your username and passphrase. If this is the first time you are logging in, enter the username and passphrase that was emailed to you by the system. You will be prompted to change your passphrase as part of your first login.

Step 3 After you log in, the OMD Director home page appears. The home page is the CDN Edge KPI metrics page. If you are logged into either an OMD Director node running in standalone mode or the primary OMD Director node in an HA installation, the title bar of the window should be solid and there will be no message showing the state, as shown below.

Note If you are missing the Media Streamer user interface options that you expect, your OMD Director is configured in Media Broadcaster only mode, which shows Media Broadcaster, Monitor, Administration, and Support menus. For details on configuring your OMD Director user interface to show the Media Streamer menus, see the *Cisco Media Streamer and Cisco Media Broadcaster Installation and Upgrade Guide*.



If you are logged into the Backup OMD Director instance, the title bar of the window will show “This OMD Director is currently running in backup mode”, as shown below.



If you are logged into an OMD Director instance that is running in the detached state, the title bar of the window will show “This OMD Director is currently running in detached mode”.

Note Each user can only have one active login at a time. If a user logs in with the same username from another system or browser, they will be prompted to either force the login, which will log off the current user session, or cancel the log in and log in as another user.

Logging out of OMD Director

To logout of OMD Director, click your username in the upper-right corner of the window and choose **Logout**.

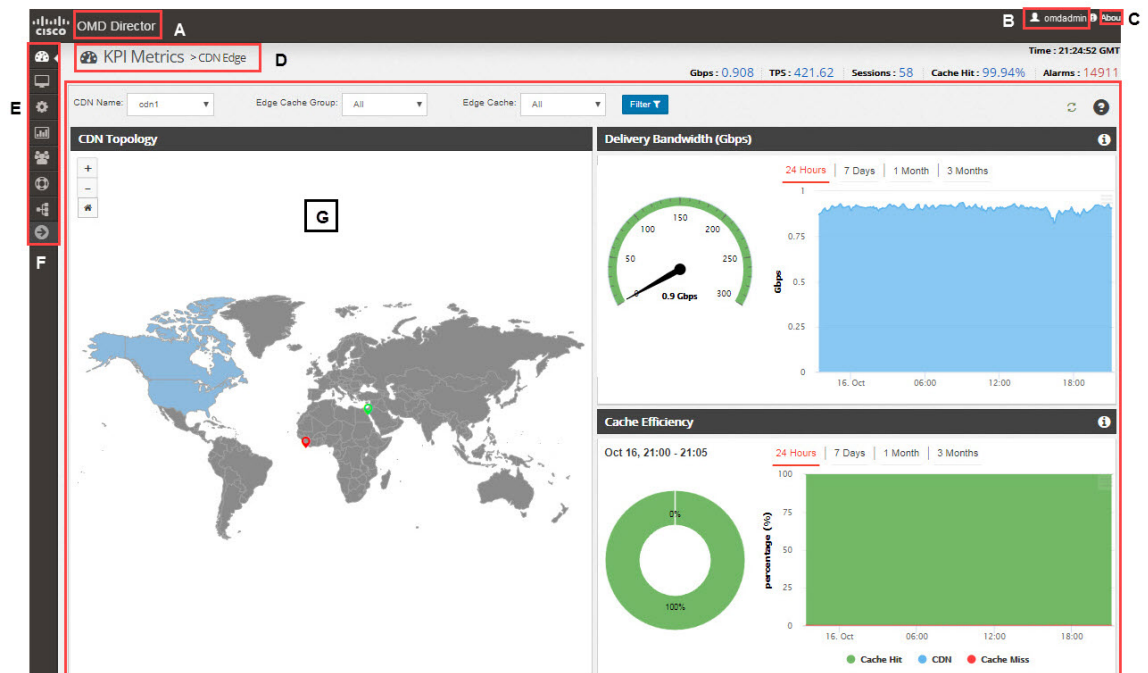
OMD Director Portal

An OMD Director portal is intended for customers with downstream Resellers or Content Providers who require access to CDN analytics. An OMD Director portal instance is a separate setup from the primary OMD Director installation, with its own user database. It provides a Director GUI that is limited to working with the CDN analytics dashboards provided by Insights. You cannot manage or view the CDN configuration from an OMD Director portal instance. An OMD Director portal is optional and does not replace the primary OMD Director Master Controller and Worker nodes.

You access the OMD Director portal GUI the same way you access the OMD Director primary GUI. However, when you are logged in to the OMD Director portal GUI, you will only be able to access the Insights, Administration, and Support menus, and the Administration menu will only be available to users with the Administrator role.

OMD Director User Interface Elements

The following diagram documents the different elements of the OMD Director user interface:



- **(A) Product Name:** Displays the name of the product (OMD Director).
- **(B) User:** Displays the user name of the currently logged in user. From this element you can logout and you can change profile settings of the current user, including the passphrase.
- **(C) About:** Clicking the about link displays the current versions of the OMD Director components. This can be helpful when reporting issues or troubleshooting problems with OMD Director.
- **(D) Current Page:** Displays the path of the current page.
- **(E) Navigation Panel:** Provides the menus to navigate to the different functions of OMD Director. The navigation panel contains the following choices:
 - KPI Metrics
 - Monitor
 - Provisioning
 - Insights
 - Administration
 - Support
 - Media Broadcaster



Note

You will only see the Media Broadcaster option if you have the Media Broadcaster installed. Also, if you are logged in to an OMD Director portal instance, you will only see the Insights, Administration (for Administrator users only), and Support menus.

- **(F) Expand Menu:** Click this icon to expand the navigation panel and display it in a menu structure. To minimize the navigation panel back to icons only, click the right arrow icon that appears.
- **(G) Main Area:** The information that is displayed in this area depends on what menu you choose.

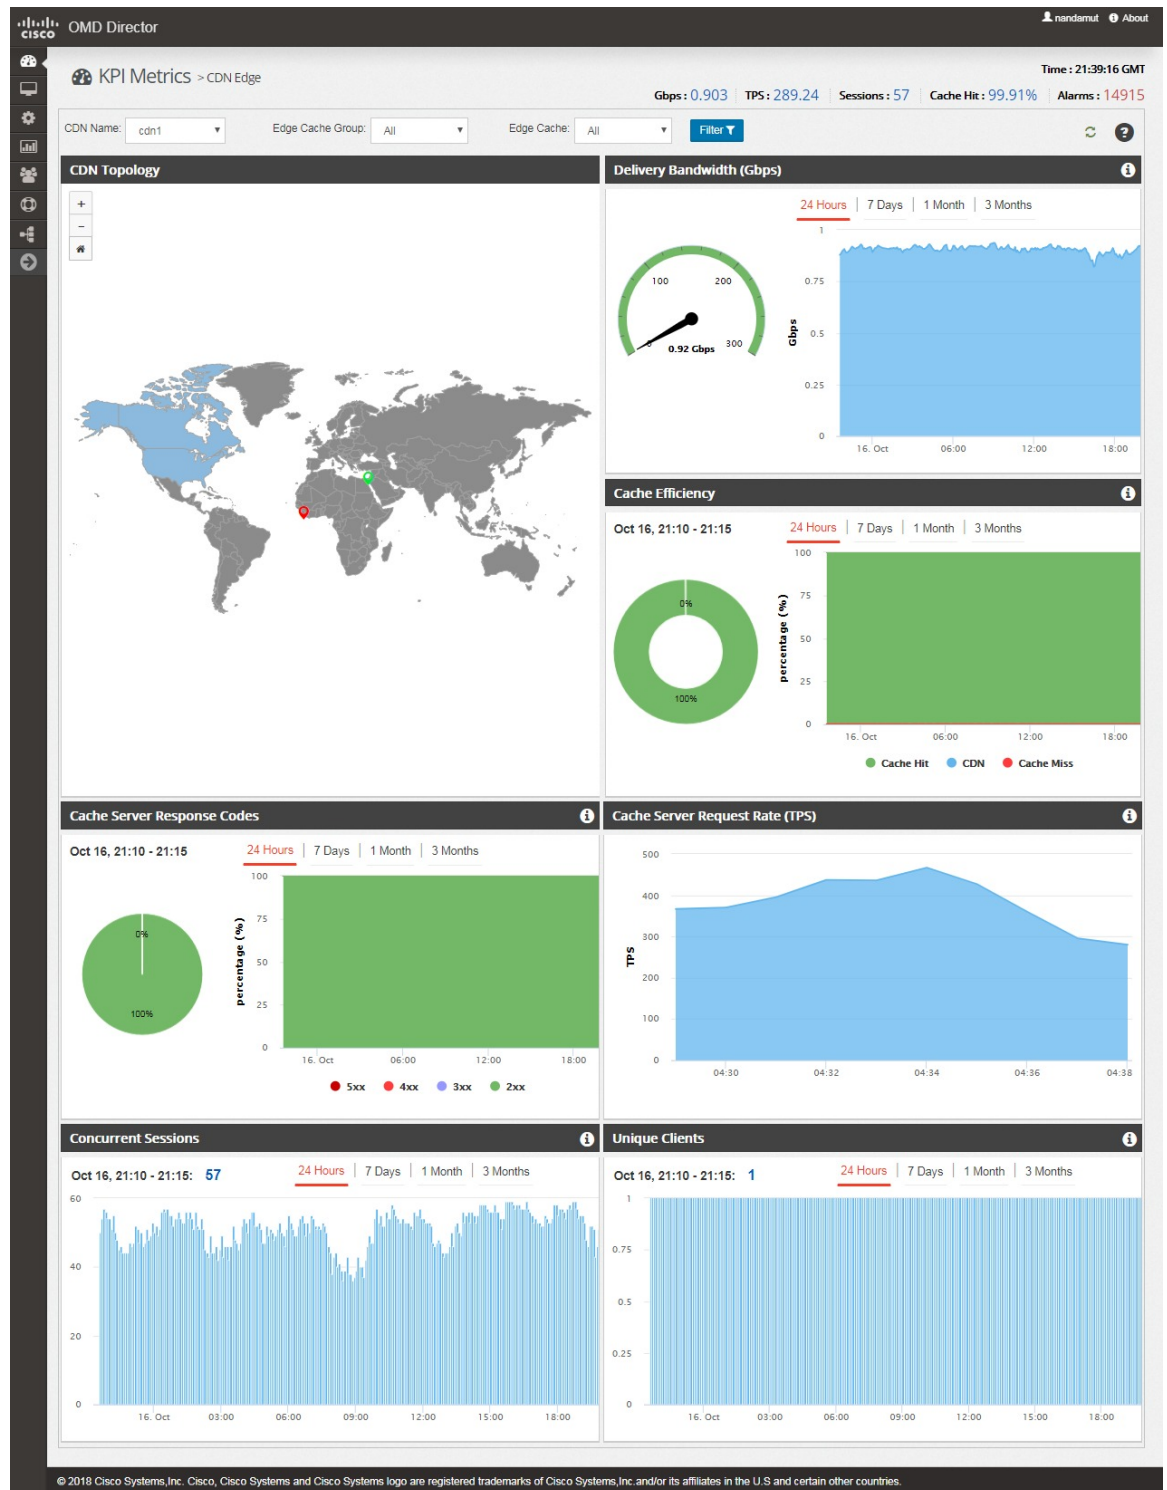
OMD Director Navigation Panel: KPI Metrics Overview

From the KPI Metrics menu you can see the key performance indicators (KPIs) for utilization of the CDN Edge caches and Mid caches. This section will provide an overview of the options available from this menu. For more detailed information on each option, see [KPI Metrics, on page 187](#).

KPI Metrics > CDN Edge

By default, the KPI Metrics > CDN Edge page provides a topology of the CDN Edge cache groups and provides statistical information across all of the Edge cache groups in the CDN. The Edge KPI Metrics window provides information about content that is delivered from Edge Caches to Clients, such as set-top boxes (STB), smartphones, and tablets.

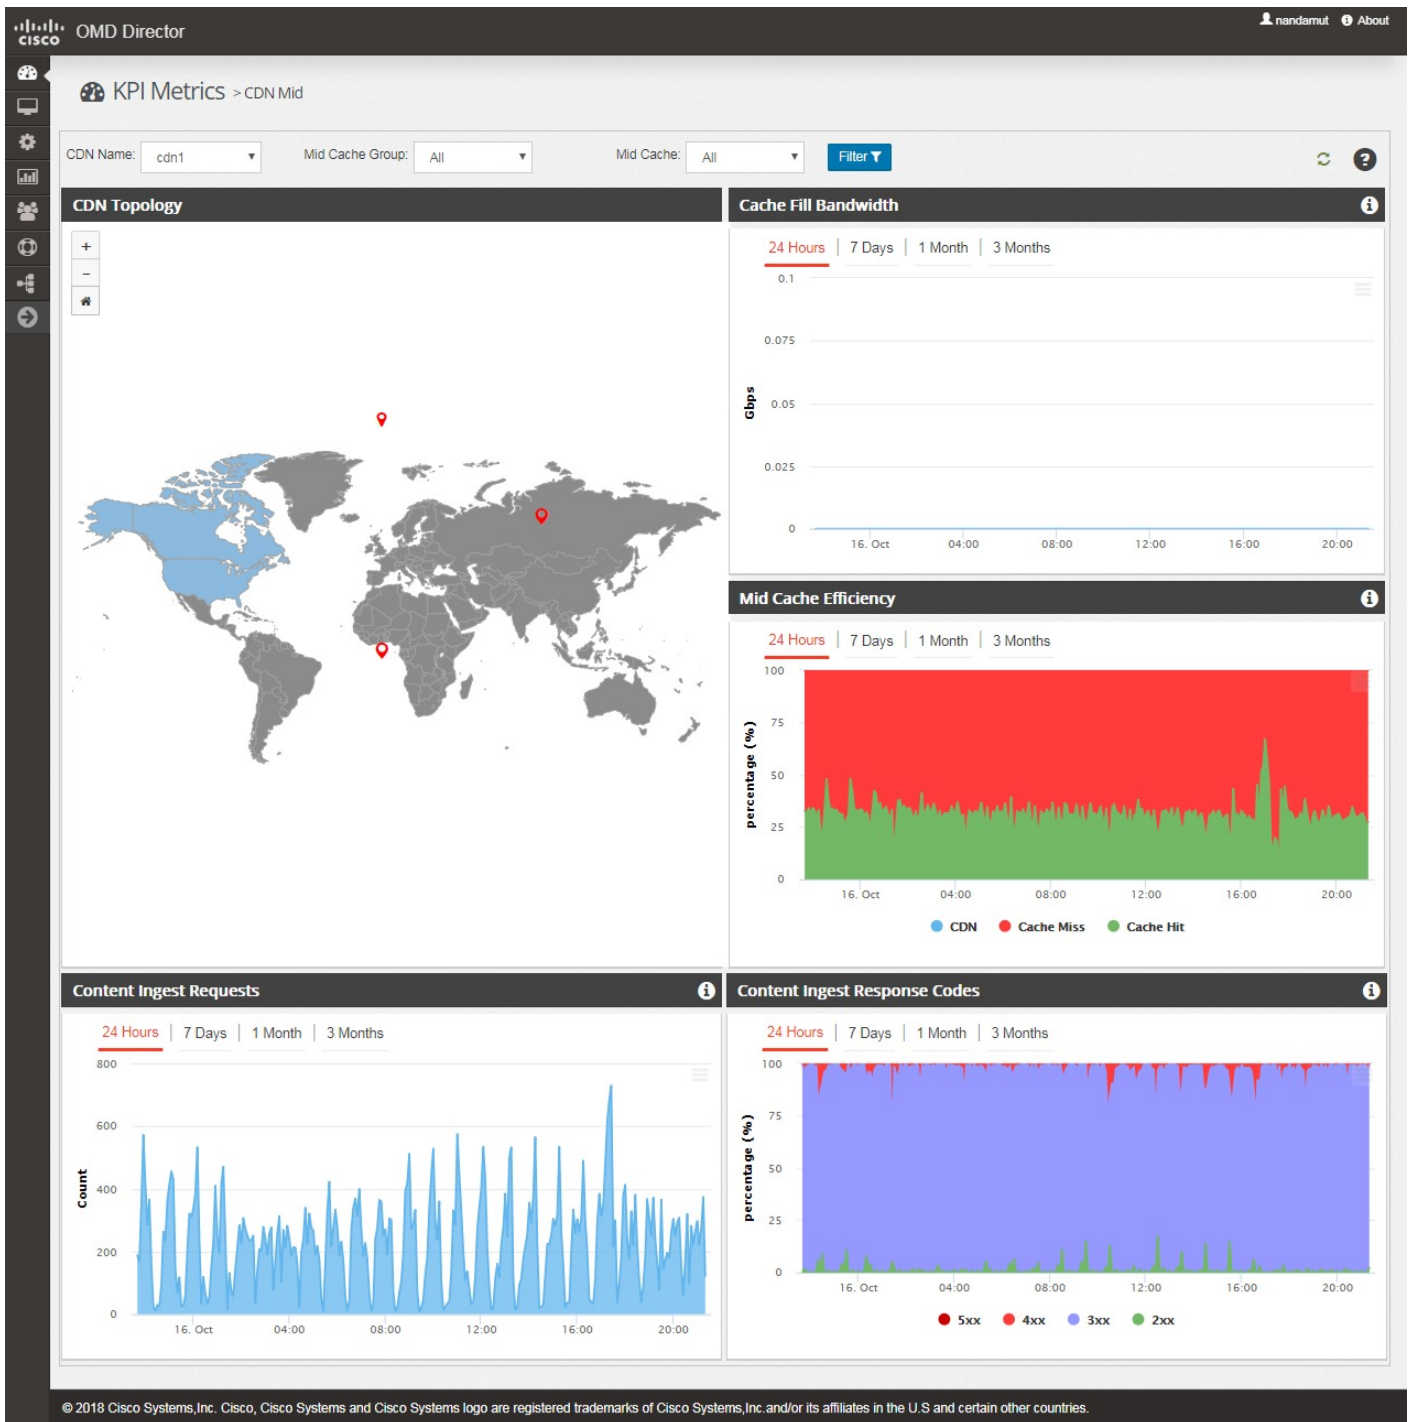
You can also filter this output by Edge Cache Group or Edge Cache. For detailed information on the information provided by the CDN Edge page and how to filter the data, see [KPI Metrics, on page 187](#).



KPI Metrics > CDN Mid

By default, the **KPI Metrics > CDN Mid** page provides a topology of the CDN Mid cache groups and displays statistical information across all of the Mid cache groups in the CDN. The CDN Mid window provides KPIs about content cache-fill into Edge Caches, and content ingest from the Origin Servers. It provides data that you can use to understand the volume of core network traffic utilization to fill edge caches.

You can also filter this output by Mid Cache Group or Mid Cache. For detailed information on the information provided by the CDN Mid page and how to filter the data, see [KPI Metrics, on page 187](#).

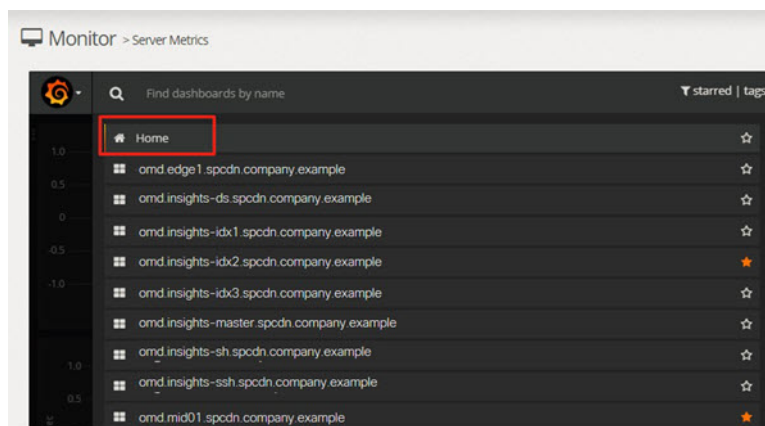


OMD Director Navigation Panel: Monitor Overview

From the Monitor menu you can see the system-level statistics for the cache servers and the generated and available CDN alarms. This section will provide an overview of the options available from this menu. For more detailed information on each option, see [CDN Monitoring, on page 201](#).

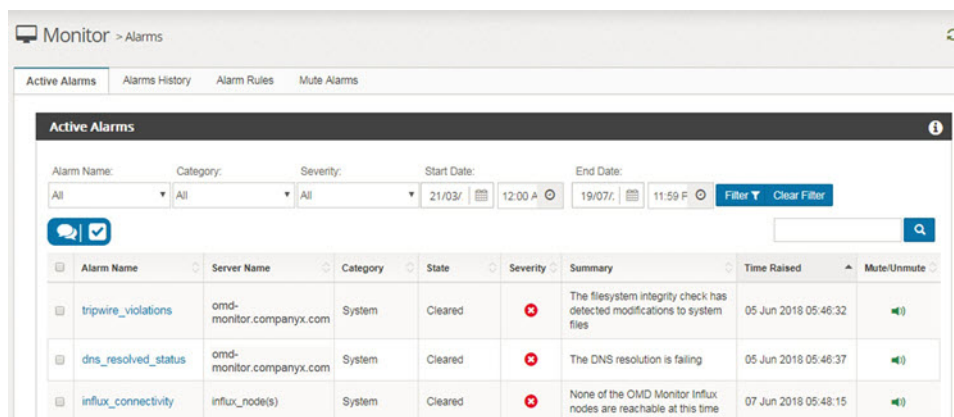
Monitor > Server Metrics

Monitor > Server Metrics launches a page that enables you to see a list of all of the physical and virtual servers used by OMD. From this list you can choose a server that will then display detailed historical server metrics about this server such as CPU utilization, Memory Utilization, Disk I/O, and other lower level server data.



Monitor > Alarms

Monitor > Alarms displays the OMD system alarms that have been raised in the CDN environment, based on thresholds that are set. From the Monitor > Alarms page you can view alarms, configure thresholds for the alarms, and mute alarms.



OMD Director Navigation Panel: Provisioning Overview

From the Provisioning menu, you can run the CDN wizard to provision the CDN, and view and manage the elements of the CDN. The Provisioning menu provides the following options:

- **CDN Wizard:** The CDN Wizard is used for the initial provisioning of the delivery services, Cache servers, Traffic Router, and Traffic Monitor of the CDN. For detailed information see [Using the CDN Wizard, on page 27](#).
- **CDN Overview:** After you have completed the CDN Wizard, CDN Overview provides a single page that displays the entire CDN, including the delivery services, cache groups, and individual servers in the CDN. For detailed information see [Validating the CDN Configuration, on page 53](#).
- **Edit CDN:** From this menu you can do the following:
 - Add and configure existing servers, cache groups, and delivery services.
 - Edit the Proximity Routing settings and Anonymous Blocking settings for the Traffic Router profiles. For more information on Proximity Routing, see [Proximity Routing, on page 159](#). For more information on Anonymous Blocking, see [Anonymous Blocking, on page 169](#).
 - Manage the CZF, NGB whitelist, and ISP Database files assigned to the Traffic Router profiles. For more information on the CZF, see [Coverage Zone File, on page 156](#). For more information on the NGB whitelist, see [NGB Whitelist, on page 176](#). For more information on the ISP Database file, see [ASN Blocking, on page 173](#).
 - Edit the Multi Site Origin (MSO) settings for the Delivery Service and Origin Server profiles. For more information on this feature, see [Multi Site Origin, on page 133](#).
 - Edit the URL Signing settings for the Delivery Service profile. For more information on this feature, see [URL Signing, on page 127](#).
 - Enable and configure the Edge Geo Blocking feature. For more information on this feature, see [Edge Geo Blocking, on page 150](#).
 - Create and manage content invalidation tasks. For more information on content invalidation policies, see [Manage Content Invalidation Policies, on page 177](#).
 - Specify Regex parameters and settings to modify the behavior of what regular expressions will match. For more information on how to configure these parameters and settings, see [Regex Remap Settings, on page 147](#).

Specify parameters and settings that you can use to modify the behavior of what regular expressions will match.
- **General Settings:** From this window you can do the following:
 - Upload an Anonymous IP database. You can use an Anonymous IP database to block client requests from Anonymous VPNs, hosting providers, public proxies, and Tor Exit nodes. For more information on using the Anonymous IP database, see [Anonymous Blocking, on page 169](#).

**Note**

A license is required to the MaxMind Anonymous IP database. For more information on obtaining this license, please contact your Cisco Account team.

- Perform basic management of profiles, including creating new profiles, changing the description of a profile, and assigning the profile to a different Media Streamer CDN.
- View Coverage Zone File (CZF) assignments and upload new CZF files. For more information on CZF, see [Manage Client Routing, on page 155](#).
- Manage Device Groups. For more information on managing Device Groups, see [Device Groups, on page 181](#).
- Upload an ISP database. You can use an ISP database to block client requests based on the AS number to which the client IP address belongs. For more information on using the ISP database, see [ASN Blocking, on page 173](#).
- Open a new window to the Traffic Ops user interface. Traffic Ops is an open source Traffic Server management application.

**Caution**

Do *not* make changes to servers, cache groups, or Delivery Services using Traffic Ops.

Health

Delivery Services

Servers

Parameters

Tools

Topology

Change Log (43)

About

Logout

UTC: 22:36:43

Search:

Profile	Host_Name	Edge Cache Group	Healthy	Admin	Connections	Mbps Out
ALL	ALL	ALL	✓	ALL	0	0
ALL	ALL	EDGE-NTN-VM	✓	ALL	0	0
ALL	ALL	EDGE-SJC-VM	✓	ALL	0	0
ALL	ALL	EDGE-Singapore	✓	ALL	0	0
ALL	ALL	EDGE-TLV	✓	ALL	0	0
ALL	ALL	EDGE_AMS	✓	ALL	0	0
ALL	ALL	EDGE_JRS	✓	ALL	0	0
ALL	ALL	EDGE_KJK	✓	ALL	0	0
ALL	ALL	EDGE-NTN	✓	ALL	0	0
ALL	ALL	EDGE_SJ	✓	ALL	0	0
ALL	ALL	EDGE_UK	✓	ALL	0	0
ALL	ALL	EDGE_USA_VA	✓	ALL	0	0
ALL	ALL	EdgeRTP	✓	ALL	0	0
ALL	ALL	MID-SJ	✓	ALL	0	0
ALL	ALL	MID-SJC-VM	✓	ALL	0	0
ALL	ALL	MID-US-VA	✓	ALL	0	0
ALL	ALL	OrphanGroup	✓	ALL	0	0

Showing 1 to 17 of 17 entries

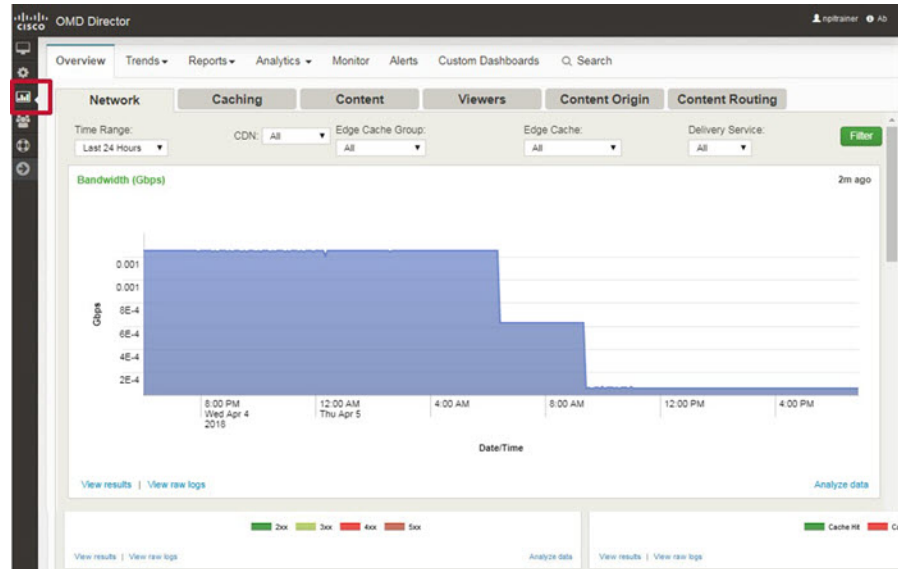
OMD Director Navigation Panel: Insights Overview

Choosing Insights from the navigation panel will launch the Media Streamer CDN Analytics powered by Splunk for big data analysis.

This page provides deep CDN insights based on mining transaction log data from the HTTP cache of each Traffic Server. The analytics information includes traffic volume, geo-distribution, content popularity, and content protocol distribution.



Note For more detailed information on the OMD Director Insights page, see [OMD Insights, on page 215](#).

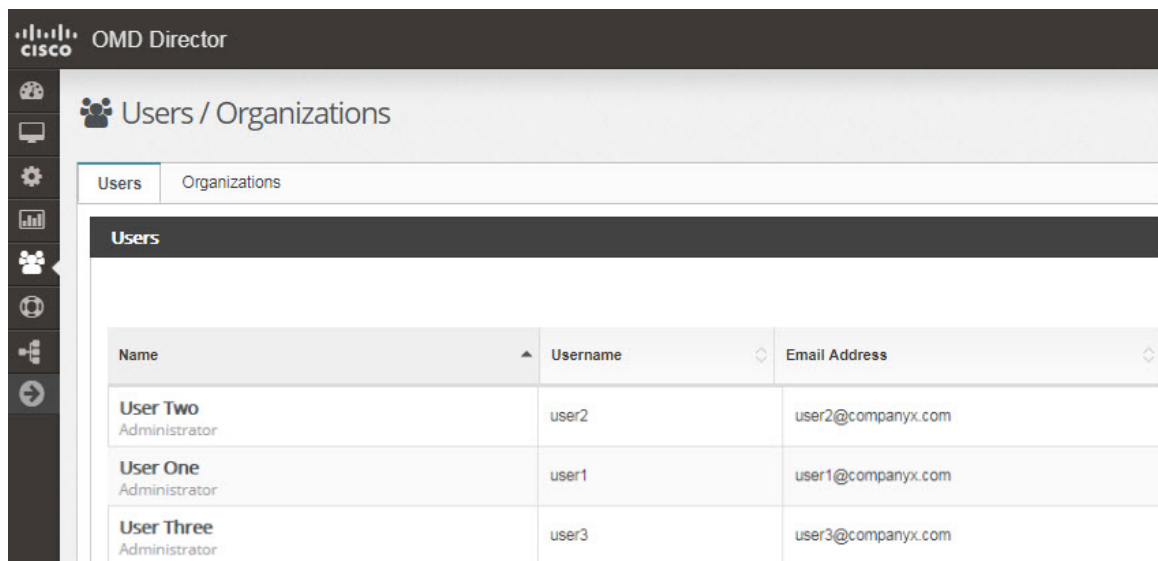


OMD Director Navigation Panel: Administration

The Administration menu enables you to see and manage the users who have access to OMD Director, including their privileges in the system. From the Administration menu you can also manage the backups of the OMD databases, including adding remote backup destination. The Administration menu also enables you to display an activity log that shows both user authentication and authorization activity, and CDN configuration activity.



Note For more information on managing the OMD Director users, see [OMD Administration, on page 235](#).



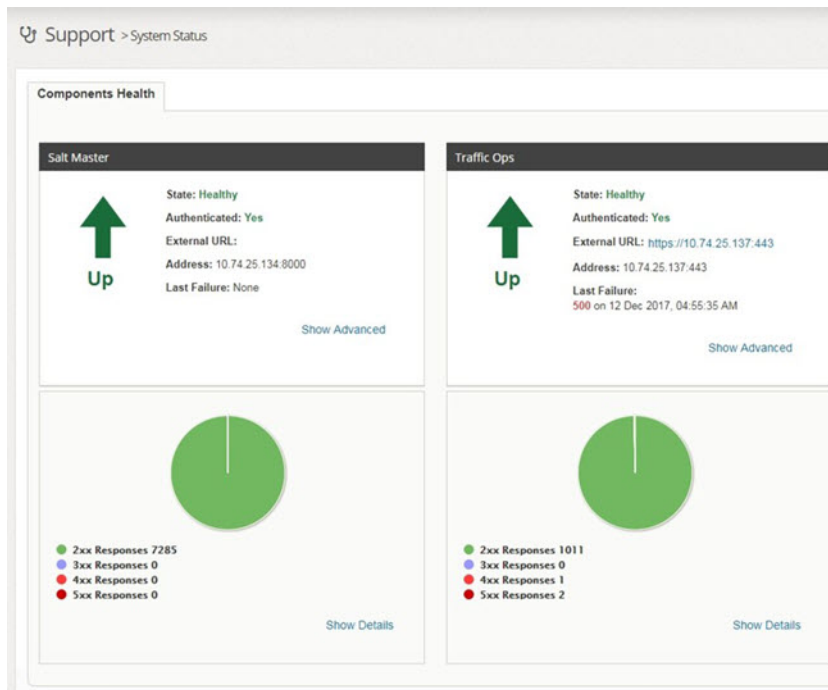
OMD Director Navigation Panel: Support

The Support menu provides the following information:

- **OMD Well Known Ports:** Provides information for the security personnel of an organization about the ports and protocols that need to be open between devices in the environment and other systems.
- **Contact:** Provides contact information for Cisco support.
- **System Status:** Displays information about the connection to the Salt Master server and the Traffic Ops server, and a historical count of the different HTTP response codes that have been received by the servers since the CDN Manager microservice was started.

When you choose **Support > System Status**, the Components Health tab is displayed. This tab displays the following information for both the Salt Master and Traffic Ops servers:

- **State:** The State will be either Healthy or Unhealthy. This state is based on an overall summary of the connection to the server based on a combination of the authentication status, last response code, and recent response history.
- **Authenticated:** This field is True if the username and passphrase that was configured can connect to server. Otherwise this field will display False.
- **Address:** Displays the Hostname or IP address, and port number that was configured for the server.
- **Last Failure:** Displays the HTTP Status Response code and timestamp of the last error on the server.
- **Show Advanced:** When you click Show Advanced, the contents of the JSON file is displayed. This JSON file is used to determine the values of the information that is displayed for the servers on the Components Health tab.
- **Response Code Table and Pie Chart:** Display a historical count of the different HTTP response codes that the server has received since the CDN Manager microservice was started.

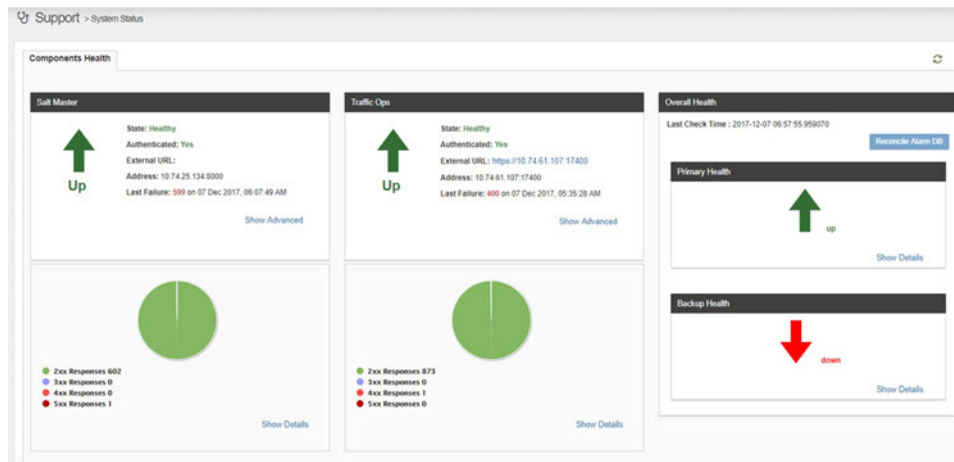


To see information about the connection to the Traffic Ops server, click the Traffic Ops Health tab. This tab displays the following information:

- **General Status:** This shows whether the server is Up or Down.
- **State:** The State will be either Healthy or Unhealthy. This state is based on an overall summary of the connection to the Traffic Ops server based on a combination of the authentication status, last response code, and recent response history.
- **Authenticated:** This field is True if the username and passphrase that was configured can connect to the Traffic Ops server. Otherwise this field will display False.
- **Address:** Displays the Hostname or IP address, and port number that was configured for the Traffic Ops server.
- **External URL:** Displays the address that cache servers can use to connect to the Traffic Ops server. This address is often the public IP address of a NAT configuration or an Openstack Floating IP address.
- **Last Failure:** Displays the HTTP Status Response code and timestamp of the last error on the Traffic Ops server.
- **Show Advanced:** When you click Show Advanced, the contents of the JSON file is displayed. This JSON file is used to determine the values of the information that is displayed on the Traffic Ops Health tab.
- **Response Code Table and Pie Chart:** Display a historical count of the different HTTP response codes that the Traffic Ops server has received since the CDN Manager microservice was started.

To refresh the information on the **Components Health** tab, click the **Refresh** button in the upper-right corner of the page.

If you are running OMD Director in an HA configuration, the Components Health window will also show the overall health of the Primary and Backup OMD Director instances.



To check the status of the individual services that run on either the Primary or the Backup, click the **Show Details** link. The following is an example:

Backup Health Details	
Name	Status
Alarms	up
MongoDB	up
Postgres	down
Notification	up

OMD Director Navigation Panel: Broadcasting Groups



Note

You will only see the Broadcasting Groups option if you have Cisco Media Broadcaster installed.

The Cisco Media Broadcaster solution transforms unicast video into multicast data that is distributed across network routers to achieve broadcast-efficient video distribution, which provides service providers a cost-effective scalable solution for streaming HTTP Live adaptive bit rate (ABR) video to in-home primary screens. Cisco Media Broadcaster integrates with an existing ABR delivery system without the need to modify the ABR clients or the video headend components, such as the encoder, packager, and content delivery network (CDN).

For more information on using OMD Director to configure broadcasting groups for Cisco Media Broadcaster, see the *Cisco Media Broadcaster User Guide*.



CHAPTER 3

Initial CDN Provisioning

After all of the prerequisites steps of the Media Streamer installation are complete, which are documented in the *Cisco Media Streamer and Cisco Media Broadcaster Installation and Upgrade Guide*, you must provision the CDN. To provision the CDN you need to do the following:

- Create the CDN
- Register servers, including Traffic Router, Traffic Monitor, and Caches, and assign them to the appropriate profiles
- Create cache groups and assign registered servers to these groups
- Create Delivery Services
- Verify the configuration



Note

If you are running OMD Director in an HA configuration, you can only make configuration changes from the Primary OMD Director instance. To perform the steps in this chapter, make sure you are logging into the Primary OMD Director instance. If you are logged into an OMD Director instance that is running in Backup mode, the header of the window will show “This OMD Director is currently running in backup mode”. If you are logged into an OMD Director running in a detached state because of an HA failure, the header will show “This OMD Director is currently running in detached mode”.

This chapter describes how to use the CDN Provisioning menu in OMD Director to perform these steps.

This chapter includes the following topics:

- [Understanding CDN Client Blocking Options, on page 23](#)
- [OMD Director CDN Deployment Steps, on page 27](#)
- [Validating the CDN Configuration, on page 53](#)

Understanding CDN Client Blocking Options

A Traffic Router handles client requests for content and determines which client requests to allow. A Traffic Router supports the following methods to determine which client requests to allow:

- **Coverage Zone File (CZF):** The CZF is a static JSON file that maps IP address ranges to cache groups. A Traffic Router can check the CZF for an IP address range that matches the requesting address to

determine if the client request is allowed. For more information and how to configure this option, see Geo Limit in the following section: [Advanced Settings: General, on page 38](#).

- **National Geoblocking:** National geoblocking (NGB) uses a geolocation database to determine what country the client request is coming from to determine if the client request is allowed. For more information and how to configure this options, see Geo Limit in the following section: [Advanced Settings: General, on page 38](#).
 - **NGB Whitelist:** An NGB whitelist is an optional whitelist that can be configured to work with the National geoblocking feature to identify addresses that NGB should not be block. For more information on the NGB whitelist and how to configure it, see [NGB Whitelist, on page 176](#).



Note CZF and NGB are referred to together as "Geo Limit".

- **Anonymous Blocking:** Anonymous Blocking uses the MaxMind Anonymous IP database to identify requests that are coming from commercial VPN services, Tor Exit Nodes, Hosting Providers, and Public Proxies. You can configure the OMD deployment to use this information to block traffic for these sources. For more information and how to configure this option, see [Anonymous Blocking, on page 169](#).



Note Anonymous Blocking is not supported by DNS Delivery Services unless Edge Geo Blocking is enabled on the Delivery Service. For more information on Edge Geo Blocking, see [Edge Geo Blocking, on page 26](#).



Note A license is required to use the MaxMind Anonymous IP database. For more information on obtaining this license, please contact your Cisco Account team.

- **Anonymous Blocking White Lists:** Anonymous Blocking white lists are configured as part of the Anonymous Blocking feature and can be used to identify addresses that should not be blocked by Anonymous Blocking. For more information and how to configure this option, see [Anonymous Blocking, on page 169](#).
- **ASN Blocking:** ASN blocking enables you to block client requests based on the autonomous system number (AS number) to which the client IP address belongs. ASN blocking uses the MaxMind ISP or the GeoLite “ASN” database to determine the AS number of an IP address. For more information and how to configure this option, see [ASN Blocking, on page 173](#).



Note ASN Blocking is not supported by DNS Delivery Services unless Edge Geo Blocking is enabled on the Delivery Service. For more information on Edge Geo Blocking, see [Edge Geo Blocking, on page 26](#).



Note A license is required to use the MaxMind ISP database. For more information on obtaining this license, please contact your Cisco Account team.

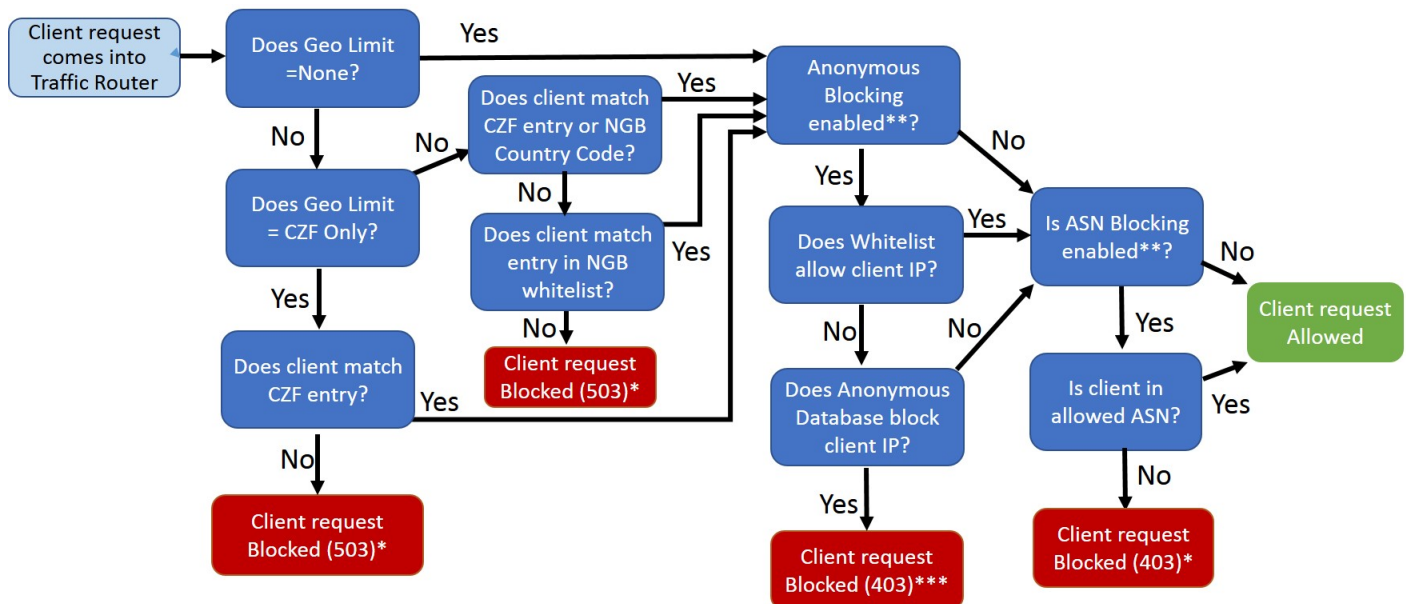
Which client blocking methods are used depends on the options that you configure for the Delivery Service. If CZF, NGB, Anonymous blocking, and ASN blocking are enabled, they are checked in the following order:

1. CZF
2. NGB
3. Anonymous blocking
4. ASN blocking



Note After the Traffic Router determines that a client request is allowed, the Traffic Router uses client routing methods to determine which cache group to use to deliver the content for that request. The Traffic Router then chooses a cache to use from that group based on cache availability, cache load, and cache content. For more information on the available client routing methods and the order in which they are used, see [Manage Client Routing, on page 155](#).

The following diagram describes the flow that a Traffic Router will use to determine whether a client request is blocked or allowed, based on which options are configured.



*Client 302 redirected to "Geo Limit Redirect URL" if configured.

**Not supported by DNS-based Delivery Services *unless* Edge Blocking is enabled on the Delivery Service

***Client 302 redirected to "Anonymous Redirect URL" if configured.

CZF and NGB are implemented by configuring the Geo Limit setting of the Delivery Service. This configuration is covered in the following sections:

- [Step 4: Create Delivery Services, on page 35](#)
- [Add a New DNS or HTTP Delivery Service, on page 92](#)
- [Edit an Existing DNS or HTTP Delivery Service, on page 124](#)

For information on configuring Anonymous Blocking, see [Configure Anonymous Blocking, on page 169](#). For more information on configuring ASN Blocking, see [ASN Blocking, on page 173](#).

Edge Geo Blocking

Without Edge Geo Blocking configured on the Delivery Service, for Delivery Services that use DNS routing the Traffic Router never sees the complete URL requested by the client or the source IP address of the client making the request. Because the Traffic Router never sees the source IP address of the client making the request, Delivery Services that use DNS routing have the following limitations:

- For Delivery Services that use CZF or NGB *and* DNS routing, the Traffic Router uses the IP address of the DNS resolver and *not* the IP address of the actual client to determine whether to allow the request.
- Anonymous blocking and ASN blocking are not supported.



Note

These limitations do *not* apply to Delivery Services that use HTTP-based routing because for HTTP-based routing requests, the Traffic Router will see the actual client IP address and full URL of the request.

The Edge Geo Blocking feature enables the Edge caches to participate in determining whether a client request should be allowed. With this feature enabled, the Edge cache passes the IP address of the actual client making the request and the requested URL to the Traffic Router to determine whether the request should be allowed. Therefore, Delivery Services that use DNS routing that have Edge Geo Blocking enabled provide the following:

- Support for anonymous blocking
- Support for ASN blocking
- For all of the blocking mechanisms (CZF, NGB, Anonymous blocking, and ASN blocking), the Traffic Router can use the IP address of the actual client to determine whether to allow the request.

Because the Edge caches now participate in determining whether a client request should be allowed, Edge Geo Blocking also ensures that clients cannot bypass blocking policies for HTTP-based or DNS-based Delivery Services by bypassing the Traffic Router.

When Edge Geo Blocking is enabled, the Traffic Router maintains all existing blocking functionality and still performs the initial check on the client request. The Edge caches have no autonomy to make blocking decisions. Edge caches can only implement the blocking instructions provided by the Traffic Router.

For additional information, including how to configure the Edge Geo Blocking feature, see [Edge Geo Blocking, on page 150](#).

OMD Director CDN Deployment Steps



Note If you are running OMD Director in an HA configuration, you can only make configuration changes from the Primary OMD Director instance. To perform the steps in this section, make sure you are logging into the Primary OMD Director instance.

Using the CDN Wizard

The CDN Wizard is used for the initial provisioning of the CDN, Delivery Services, and traffic servers (caches). You must run the CDN Wizard because this is the only way to deploy the CDN through OMD Director.



Caution Before you use this wizard, all of the prerequisite steps of the Media Streamer installation must be complete. For details, refer to the *Cisco Media Streamer and Cisco Media Broadcaster Installation and Upgrade Guide*.

To start the CDN Wizard, choose **Provisioning > CDN Wizard**. The CDN Wizard page appears.

The screenshot displays the OMD Director web interface for the CDN Wizard. The top navigation bar shows 'Provisioning > CDN Wizard' and a 'Reset' link. A progress bar at the top of the main content area indicates the sequence of six steps: 1. Create CDN (active), 2. Prepare Servers, 3. Assign Servers, 4. Create Delivery Services, 5. Review, and 6. Accept. Below the progress bar, the 'Step 1 - Create CDN' section provides instructions for creating a Content Distribution Network (CDN). It includes two input fields: 'CDN Name' with the value 'sevt-cdn' and 'CDN Domain Name' with the value 'spcdn.company.example'. At the bottom of the form, there are 'Previous' and 'Next' buttons.

There are six steps in the CDN Wizard:

Step 1: Create CDN

Step 2: Prepare Servers

Step 3: Assign Servers

Step 4: Create Delivery Services

Step 5: Review

Step 6: Accept

Step 1: Create CDN

Procedure

Step 1

In the first step of the CDN Wizards you must create the CDN. To create the CDN you must provide the following information:

- **CDN Name:** Enter a human readable name for the CDN.
- **CDN Domain Name:** Enter the CDN Domain which must match the domain for which the Traffic Router is authoritative.

Step 2

After you have confirmed the CDN Name and CDN Domain Name are correct, click **Next** or click **2 Prepare Servers** to go to Step 2.

The screenshot displays the OMD Director Provisioning > CDN Wizard interface. At the top, the breadcrumb is 'Provisioning > CDN Wizard'. Below this, a progress bar shows six steps: 1 Create CDN, 2 Prepare Servers, 3 Assign Servers, 4 Create Delivery Services, 5 Review, and 6 Accept. Step 2, 'Prepare Servers', is highlighted with a red box. Below the progress bar, the 'Step 1 - Create CDN' section is visible. It includes an 'Instructions' section with text explaining the wizard's purpose. Below the instructions are two input fields: 'CDN Name' with the value 'sevt-cdn' and 'CDN Domain Name' with the value 'spcdn.company.example'. At the bottom of the form, there are 'Previous' and 'Next' buttons. The 'Next' button is highlighted with a red box.

Step 2: Prepare Servers

In the second step of the CDN Wizard you register the existing Edge cache, Mid cache, Traffic Router, and Traffic Monitor servers that were provisioned as part of the initial Media Streamer installation. The servers that appear on the Prepare Servers page have already been provisioned with an operating system, IP address, and hostname as part of the initial installation of Media Streamer.



Note Before you can use the CDN Wizard to prepare the servers, each Media Streamer server, including the cache servers, must have the CentOS or RHEL operating system installed and configured and have the necessary prerequisite packages installed. OMD Director must also be configured to manage each cache server. For more information on these steps, refer to the *Cisco Media Streamer and Cisco Media Broadcaster Installation and Upgrade Guide*.

Follow these steps to register a server:

Procedure

Step 1

For the Edge cache, Mid cache, Traffic Router, and Traffic Monitor servers, click the **Install** icon at the end of the server row to register these servers. This will register the server in the system. If the server is a cache engine, it will also install the Media Streamer caching application on the server. While the server is being provisioned, it will show a status of “Provisioning”.

Note The Fingerprint listed is a hash of the SSH key that is used to secure the connection between the Salt Master and the Salt Minion. The fingerprint is displayed so that you can confirm that you are provisioning the correct Salt Minion.

CDN Wizard

1 Create CDN 2 Prepare Servers 3 Assign Servers 4 Create Delivery Services 5 Review 6 Accept

Step 2 - Prepare Servers

Instructions

To register a new edge or mid cache with OMD Director, please configure the hostname and networking. Then configure a SaltStack Minion to connect to OMD Director for provisioning. Verify that all cache servers in your CDN are present in this list before continuing to the next step. If a server does not appear in this list after completing the installation process, please follow the [Troubleshooting Procedure](#).

Registered Servers

Server Name	Fingerprint	Type	Status	Profile	Install	Delete
omd-edge-1	14:71:57:a2:1e:60:57:5d:d4:2d:bd:f5:25:2d:74:5a		Unregistered	-- Select Profile--		
omd-edge-2	4a:b2:92:d0:b0:2c:16:34:f3:e2:35:29:93:1b:3f:2c		Unregistered	-- Select Profile--		
omd-mid-1	37:65:4a:e9:1e:83:e1:e6:1d:0f:9d:5a:3d:20:fa:de	mid	Provision Success	-- Select Profile--		
omd-mid-2	66:ed:96:cf:01:48:e5:53:07:55:20:8f:e3:e0:ae:b8	mid	Provision Success	-- Select Profile--		

After the system is finished provisioning the server, you will see the following changes in the table:

- The Type column lists the type of server, which could be mid, edge, monitor, or router. The server type is defined during the initial installation of Media Streamer, it is not configured as part of the CDN Wizard.
- The status will say “Provision Success”.
- You will be able to choose a profile for the server.

Step 2 From the **Profile** drop-down list of a registered *cache server*, choose the Profile to assign to the server. These profiles were created as part of the initial Media Streamer installation. The following cache server profiles were created as part of the initial Media Streamer installation:

Note Be careful to assign a profile that is appropriate for the type of server (Edge or Mid) and the type of hardware you are using.

- **EDGE_CDE250_ATS_622:** This profile is intended for Edge cache servers that use the Cisco CDE 250 and will support both Live and VoD content.
- **EDGE_CDE280_ATS_622:** This profile is intended for Edge cache servers that use the Cisco CDE 280 and will support both Live and VoD content.
- **EDGE_CDE285_ATS_622:** This profile is intended for Edge cache servers that use the Cisco CDE 280 and will support both Live and VoD content.
- **EDGE_C240_M3S_LIVE_622:** This profile is intended for Edge cache servers that will support *only* Live content.
- **EDGE_ATS_622:** This is a generic profile that is intended for Edge cache servers that will support both Live and VoD content.
- **MID_CDE250_ATS_622:** This profile is intended for Mid cache servers that use the Cisco CDE 250 and will support both Live and VoD content.
- **MID_CDE280_ATS_622:** This profile is intended for Mid cache servers that use the Cisco CDE 280 and will support both Live and VoD content.
- **MID_CDE285_ATS_622:** This profile is intended for Mid cache servers that use the Cisco CDE 280 and will support both Live and VoD content.
- **MID_C240_M3S_LIVE_622:** This profile is intended for Mid cache servers that will support *only* Live content.
- **MID_ATS_622:** This is a generic profile that is intended for Mid cache servers that will support both Live and VoD content.

If you do not see an appropriate profile for your cache server, contact your Cisco Account team.

Step 3 From the **Profile** drop-down list of a registered *Traffic Router*, choose the appropriate profile for the Traffic Router. The initial profile that was created is named TRAFFIC_ROUTER.

Step 4 From the **Profile** drop-down list of a registered *Traffic Monitor*, choose the appropriate profile for the Traffic Monitor server. This initial profile that was created is named TRAFFIC_MONITOR.

CDN Wizard

1 Create CDN 2 Prepare Servers 3 Assign Servers 4 Create Delivery Services 5 Review 6 Accept

Step 2 - Prepare Servers

Instructions

To register a new edge or mid cache with OMD Director, please configure the hostname and networking. Then configure a SaltStack Minion to connect to OMD Director for provisioning. Verify that all cache servers in your CDN are present in this list before continuing to the next step. If a server does not appear in this list after completing the installation process, please follow the [Troubleshooting Procedure](#).

Registered Servers

Server Name	Fingerprint	Type	Status	Profile	Install	Delete
omd-edge-1	14:71:57:a2:1e:60:57:5d:d4:2d:bd:f5:25:2d:74:5a		Unregistered	-- Select Profile--		
omd-edge-2	4a:b2:92:d0:b0:2c:16:34:f3:e2:35:29:93:1b:3f:2c		Unregistered	-- Select Profile--		
omd-mid-1	37:65:4a:e9:1e:83:e1:e8:1d:0f:9d:5a:3d:20:fa:de	mid	Provision Success	-- Select Profile--		
omd-mid-2	66:ed:96:cf:01:48:e5:53:07:55:20:8f:e3:e0:ae:b8	mid	Provision Success	-- Select Profile--		

Step 5 After you have finished registering all of the servers, check to make sure that the correct profile is still selected for each server, based on the server type. If the correct profile is not selected, reselect the profile before proceeding to Step 3 of the CDN Wizard.

Step 6 Click **Next** or click **3 Assign Servers** to go to Step 3.

Step 3: Assign Servers

After the servers have been successfully provisioned, you need to create cache groups and assign the registered servers to the appropriate cache groups. Part of this configuration is to assign the Traffic Monitor and the Traffic Router to the CtrlPlaneGroup cache group.

Assign Traffic Monitor and Traffic Router

To assign the Traffic Monitor and the Traffic Router to the CtrlPlaneGroup, perform the following steps:

Procedure


Step 1 From the **Select Cache Group** drop-down list, choose **CtrlPlaneGroup**.

Step 2 From the **Registered Servers** list, press and hold the **Ctrl** key while you click both the **Traffic Monitor** and **Traffic Router** servers to choose both servers. Click the right arrow button to move the servers to the Assigned Servers list.

Step 3 - Assign Servers

Instructions

On this page, create new cache groups and then assign individual servers to the cache groups. Cache groups are a logical grouping of servers that provide content to a subset of clients (i.e. a cache group to serve a certain geography). This CDN will have two tiers, an edge tier and a mid-edge cache group with a parent of one mid cache group.

Select Cache Group: 

Registered Servers

Assigned Servers

omd-monitor
omd-tr

»

➤

⬅

⬅

After the Traffic Monitor and Traffic Router servers have been successfully added to the CtrlPlaneGroup, you need to create the cache groups and assign the cache servers to their appropriate group. OMD uses a two-tiered CDN Traffic Server (cache) hierarchy:

- **Edge caches:** Provide edge caching, content streaming, and download to subscriber IP devices. Traffic routers redirect client requests to Edge caches based on geolocation, server availability, server load, and server cache content to provide efficient system-wide load balancing. Edge caches are organized into cache groups. You configure each Edge cache group a primary Mid parent cache group and optionally a secondary Mid parent cache group for failover.
- **Mid caches:** Provide content ingest and storage functionality. When an Edge cache does not contain the content requested by the client, the Edge cache will proxy the request to a Mid cache server, based on the parent cache group assigned to the Edge cache. If the Mid cache server does not contain the content, it is responsible for fetching the content from the Origin Server. Mid caches are also organized into cache groups. Mid cache groups may serve (be a parent to) multiple Edge cache groups.

Create Mid Cache Groups

Because Edge cache groups are assigned Mid cache groups as parents, you must create the Mid cache groups first. Perform the following steps to create Mid cache groups:

Procedure

- Step 1** Next to the **Select Cache Group** drop-down list, click the + icon to create a new cache group.

Step 3 - Assign Servers

Instructions

On this page, create new cache groups and then assign individual servers to the cache groups. Cache groups are a logical grouping of servers that provide redundancy for each other. Cache groups are also typically associated with a subset of clients (i.e. a cache group to serve a certain geography). This CDN will have two tiers, an edge tier and a mid-tier, so you must create at least one edge cache group with a parent of one mid cache group.

- Step 2** In the Add New Cache Group Window that appears, enter the following information for the new group:
- In the **Name** field, enter a descriptive name for the group. This is the name that appears in the OMD Director drop-down menus.
- Note** The name should begin with MID (in any upper case or lower case combination) so you can easily identify what type of group it is.
- In the **Short Name** field, enter an additional descriptive name. This name is not currently used in OMD Director, but it is a required field. The short name can be the same as the Name.
 - The **Geo Magnetic Latitude** and **Geo Magnetic Longitude** fields define the geolocation of the cache group. The geolocation is used by the CDN system to help redirect user clients to the most optimal cache. Enter the latitude and longitude that geolocation should use for this group.
 - From the **Type** drop-down list, choose **MID_LOC**.
- Step 3** Click **Add** to add the group.

Create Edge Cache Groups

After you have created the Mid cache groups, you need to create the Edge cache groups. Perform the following steps to create the Edge cache groups:

Procedure

- Step 1** Next to the **Select Cache Group** drop-down list, click the + icon to create a new cache group.
- In the **Name** field, enter a descriptive name for the group. This is the name that appears in the OMD Director drop-down menus.
- Note** The name should begin with EDGE (in any upper case or lower case combination) so you can easily identify what type of group it is.

- b) In the **Short Name** field, enter an additional descriptive name. This name is not currently used in OMD Director, but it is a required field. The Short Name can be the same as the Name.
- c) The **Geo Magnetic Latitude** and **Geo Magnetic Longitude** fields define the geolocation of the cache group. The geolocation is used by the CDN system to help redirect user clients to the most optimal cache. Enter the latitude and longitude that geolocation should use for this group.

Note For more information on how the Traffic Router determines which cache group should service a request, see [Client Routing Overview, on page 155](#).

- d) From the **Type** drop-down list, choose **EDGE_LOC**.
- e) From the **Parent Cache Group** drop-down list, choose the Mid cache group to use as a parent. Optionally you can choose a Mid cache group to use as the Secondary Parent Cache Group for failover. When the Mid cache server from the primary parent cache group that contains the content is unavailable, the secondary Mid cache group is used.

Step 2 Click **Add** to add the group.

Note For Edge cache groups, you can configure backup Edge cache groups to use when there are no caches available in the original Edge cache group selected by the Traffic Router for the client request. Backup edge cache groups cannot be configured from the CDN Wizard. After you have completed the CDN Wizard, you can edit the Edge cache group to configure backup edge cache groups. For more information on backup edge cache groups see [Backup Edge Cache Groups, on page 78](#).

Assign Registered Cache Servers

After the groups are created, you must assign the registered cache servers to their appropriate cache group. Perform the following steps to accomplish this task:

Procedure

Step 1 From the **Select Cache Group** drop-down list, choose the group to which you want to assign a cache server.

Step 2 From the **Registered Servers** list, choose the cache servers that you want to assign to the group and click the right arrow button to move the servers to the Assigned Servers list. To choose more than one server, press the Ctrl key while you click the servers.

Note The type of group that you choose will determine which Registered Servers are available to assign to the group. (You can only assign servers that have a Type of Edge to Edge cache groups and you can only assign servers that have a Type of Mid to Mid cache groups.)


Step 3 Repeat Step 2 until you have assigned all of the registered servers. To remove a server from a group, choose it in the **Assigned Servers** list and click the left arrow button.

Note If you want to assign all of the servers listed in the **Registered Servers** list to the selected group, you can click the double right arrow button.

Step 3 - Assign Servers

Instructions

On this page, create new cache groups and then assign individual servers to the cache groups. Cache groups are a logical grouping of servers that other. Cache groups are also typically associated with a subset of clients (i.e. a cache group to serve a certain geography). This CDN will have two mid-tier, so you must create at least one edge cache group with a parent of one mid cache group.

Select Cache Group: 

Registered Servers

- omd-edge-01

Assigned Servers

»

>

<

«

Step 4 When you are finished assigning all of the registered servers to their groups, click **Next** or click **4 Create Delivery Services**.

Step 4: Create Delivery Services

The next step in the CDN Wizard is to create a Delivery Service. A Delivery Service is a CDN representation of an Origin Server. It defines the URL that represents the Origin Server and which cache groups can serve content from that server, as well as policies for serving content to the clients.

You must create at least one Delivery Service.

Procedure

- Step 1** In the Step 4 - Create Delivery Services window, enter a unique name for the Delivery Service in the **Delivery Service Name** field.
- Step 2** **Display Name** field, enter a display name for the Delivery Service. This can be the same as the Delivery Service Name or you can enter a more descriptive display name.
- Step 3** In the **Long Description** field, enter a description for the Delivery Service.
- Step 4** In the **Routing Type** area, choose one of the following routing types:

Note You cannot change the value of the Routing Type after the Delivery Service is created.

- **DNS:** If you choose DNS, the Traffic Router responds with an A record in response to the client request, not an HTTP 302 redirect. This A record contains a list of IP addresses for the caches in the cache group that the Traffic Router selects. Which cache the client uses from this list is based on the Proxy DNS of the client. For a DNS Delivery Service, by default the client receives a URL with **edge** prepended to the CDN domain name, for example edge.omdcdn.example.com. To use a custom FQDN for the delivery service, add a sub domain to the delivery service with the custom FQDN. See the Sub Domain setting below.

Which IP address the Traffic Router uses to select a cache group is based on the following:

- If the Enhanced DNS Request Routing (ECS) feature is enabled in the Traffic Router profile and the ECS option is present in the DNS query, the Traffic Router will use the client subnet that is in the ECS option to select the cache group. If there are multiple ECS options in the Optional Record, the one with the longest IP prefix is used.
- If the Enhanced DNS Request Routing (ECS) feature is *not* enabled in the Traffic Router profile or if the ECS option is not present in the DNS query, the Traffic Router will use the IP address of the DNS resolver that is making the request to select the cache group.

For information on configuring the Enhanced DNS Request Routing feature, see [Configuring Enhanced DNS Request Routing, on page 271](#).

- **HTTP:** If you choose HTTP, the Traffic Router uses an HTTP 302 Redirect to control which cache the client will use. For an HTTP Delivery Service, by default the client receives a URL with `tr` prepended to the CDN domain name, for example `tr.omdcdn.example.com`. To use a custom FQDN for the delivery service, add a sub domain to the delivery service with the custom FQDN. See the Sub Domain setting below.

Step 5

In the **Content Type** area, choose one of the following content types:

Note You cannot change the value of the Content Type after the Delivery Service is created.

- **Live:** Live content is not committed to disk storage. Live video segments are only buffered into the RAM disk.
- **VOD:** VOD content is cached in disk storage and the most popular objects are stored in RAM cache.

Note The Content Scope for VOD content is always National and you cannot change this.

Step 6

In the **Content Scope** area, choose one of the following content scopes for live content:

Note You cannot change the value of the Content Scope after the Delivery Service is created.

- **National:** A client request will go to an Edge cache. If the content is not available on the Edge cache, the Edge cache will request the content from a Mid cache in the parent Mid cache group. If none of the Mid caches in the parent Mid cache group are available, the Edge cache requests the content from a Mid cache in the secondary parent Mid cache group. If the Mid cache does not contain the requested content, it will request the content from the Origin Server.

If none of the Mid caches are available in either the parent or the secondary parent Mid cache group, the "Edge Cache Retrieval from Origin" setting of the Delivery Service determines whether the Edge cache can request the content directly from the Origin Server.

- **Regional:** Content that has a Regional scope allows an Edge cache to request the content directly from the Origin Server without first requesting the content from the parent Mid cache group. This prevents Mid tier caches from having to cache regional content that it might never use.

Note The Content Scope for VOD content is always National and you cannot change this.

Edge Cache Retrieval from Origin: The Edge Cache Retrieval from Origin setting only applies to Delivery Services that have a National content scope. When a Delivery Service has a National content scope, if the content the client is requesting is not available on the Edge cache and none of the Mid caches are available in either the parent or the secondary parent Mid cache group, the value of "Edge Cache Retrieval from Origin" setting determines whether the Edge cache can request the content directly from the Origin Server:

Step 7 From the **Protocol** drop-down list, choose HTTP. For the Delivery Services that you create using the CDN Wizard, only HTTP is available. To add a Delivery Service that uses HTTPS, HTTP and HTTPS, or HTTP to HTTPS, you must go to Provisioning > Edit CDN to add the Delivery Service. For details on adding a Delivery Service from the Provisioning > Edit CDN menu, see the [Add a New DNS or HTTP Delivery Service](#) section.

Step 8 In the **Origin Server Base URL** field, enter the URL of the actual Origin Server that hosts the content for this Delivery Service.

Note To support MSO and device groups, every Delivery Service *within* a CDN must have a unique Origin Server Base URL (also referred to as the OFQDN), with the following exception: If two Delivery Services use the exact same MSO settings and use the exact same Edge caches, Mid caches, Origin Servers, and Device Groups, then two Delivery Services can use the same OFQDN.

Step 9 The **Customer** field is for backwards compatibility only. For OMD Director 3.11 and later, use the **Content Provider** field instead.

Note If there is a value entered in the **Content Provider** field and the **Customer** field, the **Customer** field is ignored.

Step 10 From the **Content Provider** drop-down list, optionally choose the Content Provider organization to associate with this Delivery Service. This is an optional field, however, the Content Provider field plays an important role in controlling who can view the OMD Insights analytics information for this Delivery Service.

- **Content Provider Viewers and Content Provider Admins:** OMD Director users that have the role of Content Provider Viewer or Content Provider Admin can only view OMD Insights information for Delivery Services that are assigned a Content Provider that matches the Content Provider organization they are assigned.
- **Reseller Viewers and Reseller Admins:** OMD Director users that have the role of Reseller Viewer or Reseller Admin can only view OMD Insights information for Delivery Services that are assigned a Content Provider that is assigned to the Reseller organization they are assigned.

For example, if user jsmith has the role of Content Provider Admin and has been assigned the Organization of CompanyX, and a Delivery Service named VOD-SJ has been assigned a Content Provider of CompanyY, jsmith *cannot* view any of the OMD Insights information for the VOD-SJ Delivery Service.

Note The Content Provider field only affects users that have one of the Content Provider or Reseller roles. It does not limit or restrict users that have been assigned any other role.

Step 11 From the **Active** drop-down list, choose Yes to enable the Delivery Service.

Step 12 In the **Assign to** field, choose the Edge caches that should serve content for this Origin Server. You do not need to configure any Mid cache groups because they are inherently assigned through the Edge cache parent relationship.

Step 13 In the **Sub Domain** field, you can create a sub domain in two different ways:

- **Enter only the sub domain:** Enter only the sub domain portion of the FQDN. This subdomain is prepended to the CDN domain and either "tr" (for an HTTP Delivery Service) or "edge" (for DNS Delivery Service) is the hostname used to create the URL for client requests to this delivery service. For example, if the CDN domain name is spcdn.company.example and you enter a sub domain of on-site, the client will use a URL for this sub domain of tr.on-site.spcdn.company.example for an HTTP Delivery Service and a URL of edge.on-site.spcdn.company.example for a DNS Delivery Service.

The following is an example:

Domains for this Delivery Service:

Sub Domain * Sample Client URL

on-site http://tr.on-site.spcdn.company.example

- **Enter a full FQDN:** A custom FQDN enables you to use a different hostname than the default hostname of “tr” for an HTTP Delivery Service or “edge” for the DNS Delivery Service. For this option you enter the full FQDN: hostname, sub domain you want to create, and domain of the CDN. for example `srv1.on-site.spcdn.company.example`.

Note You can only add one custom FQDN to a delivery service.

The following is an example:

Domains for this Delivery Service:

Sub Domain * Sample Client URL

srv1.on-site.spcdn.company.example http://srv1.on-site.spcdn.company.example

- Step 14** To configure the advanced settings for the Delivery Service, click the **Advanced Settings** link below the Assign to field. The Additional Info window appears. See the following sections for information on the fields in the Additional Info window.
- Step 15** When you are finished entering the information for the Delivery Service, click **Save**. To confirm that the Delivery Service is saved, click the **View Delivery Services** link.
- Step 16** Repeat Step 1 to Step 15 for any additional Delivery Services that you need to create. When you are finished entering all of the Delivery Services, click **Next** or click **5 Review**.

Delivery Service Advanced Settings

The Additional Info window is divided into four sections: General, Content Preposition, Header Rewrite Configurations, and URL Signing. Use the following information to configure the advanced settings available in these sections.

Advanced Settings: General

- **Query String Handling:** Query strings are the set of key/value pairs that occur after the ? in a URL. The Query String Handling setting enables you to control how requests that contain query strings are handled by the CDN. The options for this setting are:
 - **0 - use qstring in cache key, and pass up:** With this setting, the query string is saved as part of the URL cache key on the cache and will be used to identify content. Therefore, two URLs that are the same except for the query string will be saved as separate URL cache keys. In addition, the query string is passed up to either the Mid cache or Origin Server. This option is good to use when the content that is returned by the Origin Server depends on the query string. This is the default setting.
 - **1 - ignore in cache key, and pass up:** With this setting, the query string is not saved as part of the URL cache key. Therefore, two URLs that are the same except for the query string will match the same URL cache key. In addition, the query string is passed up to either the Mid cache or Origin

Server. This option is good to use when the content that is returned does not depend on the query string, but the Origin Server still requires the query string for other purposes, such as authentication or logging.

- **2 - drop at edge:** With this setting, the query string is not saved as part of the URL cache key. Therefore, two URLs that are the same except for the query string will match the same URL cache key. With this setting, the query string is not passed up to either the Mid cache or Origin Server.
- **Regex Remap Expression:** This setting enables you to configure a regular expression remap rule that modifies the URL that the client is requesting. The base remap rule modifies the original client request for retrieving content from the origin listed in the Base Origin Server URL field. Placing a rule in this field can override and replace the request URL. If the regular expression does not match the request URL, the Base Origin Server URL is used.



Note This feature cannot be used when the Query String Handling setting is set to “2 - drop at edge”.

The following is the syntax for this field:

regex_expression_fromURLpath remap_toURL options

- *regex_expression_fromURLpath:* A regular expression that is evaluated to determine which client requests to remap. Any client requests that match this regular expression are remapped. By default, the regular expression only matches against the URL path and query string. The path will always start with a “/”.
- *remap_toURL:* The URL to which you want to remap the matching client requests. Various substitution strings are allowed in the remap_toURL parameter during evaluation. For a list of substitution strings that can be used in this parameter, please refer to the open source documentation at https://docs.trafficserver.apache.org/en/5.3.x/reference/plugins/regex_remap.en.html.
- *options:* Additional options that you want to apply to the remap rule. For a list of available options, please refer to the open source documentation at https://docs.trafficserver.apache.org/en/5.3.x/reference/plugins/regex_remap.en.html.

For example:

```
^/(vod.*)/more http://www.vod.example/$h/$0/$1
```



Note By default, only the path and query string of the URL are provided for the regular expressions to match. The Regex Remap Settings tab of the Delivery Service profile enables you to specify parameters and settings that you can use to modify the behavior of what regular expressions will match. For more information on how to configure the Regex Remap parameters, see [Regex Remap Settings, on page 147](#).

- **Geo Limit:** The Geo Limit setting enables you to configure the Traffic Router to block client requests based on the coverage zone file (CZF) and the geolocation of the client IP address, depending on the option that you choose. The options for this setting are:

- **None:** When Geo Limit is set to None, geo blocking is *not* enabled and the Traffic Router will check the Anonymous Blocking configuration to determine whether to allow the request. If Anonymous Blocking is not enabled, then ASN blocking is checked to determine whether to allow the request. If ASN Blocking is not enabled, the Traffic Router will allow all client requests. None is the default setting for Geo Limit.

For more information on Anonymous Blocking, see [Anonymous Blocking, on page 169](#). For more information on ASN Blocking, see [ASN Blocking, on page 173](#).

- **CZF only:** With this setting the Traffic Router will only allow requests from clients whose IP addresses match an entry in the CZF file. If the client IP address does *not* match an entry in the CZF file, then the Traffic Router will reject the client request.



Note The CZF file is discussed in more detail in the "Manage Client Routing" chapter. Also an example of the CZF file can be found in the "Example CZF File" appendix.

- **CZF + CountryCode(s):** With this setting the Traffic Router will take the following steps:



Note If you are setting Geo Limit to "CZF + CountryCode(s)", confirm that the MaxMind geolocation database has been configured in the Media Streamer deployment. For more information on configuring Media Streamer to use the MaxMind database, refer to the "Configure the MaxMind Database for Geolocation" appendix in the *Cisco Media Streamer and Cisco Media Broadcaster Installation and Upgrade Guide*.

1. First the Traffic Router checks the CZF file. If the client IP address matches an entry in the CZF file, the Traffic Router will then check the client request against the Anonymous Blocking configuration to determine whether the request is allowed.
2. If the client IP address does *not* match an entry in the CZF file, then the Traffic Router looks up the client IP address in a geolocation database to determine which country the client IP address is in. If the client IP address is in one of the countries listed in the Geo Limit Country Codes field, the Traffic Router will then check the client request against the Anonymous Blocking configuration to determine whether the request is allowed.
3. If the client IP address is not in the CZF file and the client IP address is not in one of the countries listed in the Geo Limit Country Codes field, the Traffic Router checks the NGB whitelist if configured. If the client IP address matches an entry in the NGB whitelist, Anonymous Blocking is checked next to determine whether the request is allowed. If there is no NGB whitelist configured or if the client IP address does not match an entry in the NGB whitelist, the Traffic Router will reject the client request.

- **Geo Limit Country Codes:** Choose the country codes for the countries you want to permit client requests from. This field is used when Geo Limit is set to **CZF + CountryCode(s)**.
- **Geo Limit Redirect URL:** If the Geo Limit field is set to either "CZF only" or "CZF + CountryCode(s)", you can enter a URL in this field that the Traffic Router will redirect the client to if their request is blocked. If you leave this field blank, the client will receive a "503 Service Unavailable" response.



Note You will only see this field if the Geo Limit is set to either "CZF only" or "CZF + CountryCode(s)".

- **Bypass FQDN:** The Bypass FQDN is used in the following situations:

- **Thresholds exceeded:** If either the "Maximum Bits per Second Allowed Globally" or the "Maximum Transaction Allowed Globally" limits are reached, the Traffic Router will send any overflow traffic to the FQDN listed in the Bypass FQDN field. To redirect the requests to the Origin Server, you would enter the FQDN of the Origin Server. By configuring Bypass FQDN, the Maximum Bits per Second Allowed Globally, and the Maximum Transaction Allowed Globally settings, you can essentially place a quota on how much traffic and how many computing resources the Delivery Service can use.



Note You do not have to configure both the Maximum Bits per Second Allowed Globally and the Maximum Transaction Allowed Globally settings. However, you need to configure at least one of them for the Bypass FQDN setting to have any affect.

- **No caches available:** If there are no caches available when the Traffic Router tries to route the client request, the Traffic Router will redirect the client request to the FQDN listed in the Bypass FQDN field. To redirect the requests to the Origin Server, you would enter the FQDN of the Origin Server. This essentially works like a "Last Resort Routing".
- **Maximum Bits per Second Allowed Globally:** This setting is used in conjunction with the Bypass FQDN setting. The value that you enter in this field determines the maximum bits per second (bps) this Delivery Service can serve across all Edge caches. When this limit is reached, the Traffic Router diverts the overflow traffic to the Bypass FQDN destination.



Note For this setting to have any effect, you must also configure the Bypass FQDN setting.

- **Maximum Transaction Allowed Globally:** This setting is used in conjunction with the Bypass FQDN setting. The value that you enter in this field determines the maximum transactions per second that this Delivery Service can serve across all Edge caches. When this limit is reached, the Traffic Router diverts the overflow traffic to the Bypass FQDN destination.



Note For this setting to have any effect, you must also configure the Bypass FQDN setting.

- **Traffic Router Log Request Headers:** This field enables you to specify headers that you would like to include in the Traffic Router access log entries. To specify which headers to include, enter a list of header keys separated by __RETURN__. For example, "Header1-Name __RETURN__ Header2-Name". You

can add additional headers by separating them with "__RETURN__". These headers will appear after the "rh=" token. If the listed header name is not present, it will not be logged.

- **Initial Dispersion:** The Initial Dispersion setting determines how many copies of requested content are stored in the cache group. By default, the Initial Dispersion setting is 1, which means that only a single copy of the requested content is cached in the cache group. Therefore, by default only one cache server will have a copy of the content. If the cache server that contains the content goes offline or is overloaded, the content is moved to another cache server in the cache group. This optimizes the amount of storage that is available for caching content.

However, if you have a piece of content that you know is going to be very popular, you may want OMD to cache more than one copy of the content. These copies would be created at the time of the initial request and the system would not have to wait for load or cache health settings to be reached before additional copies of the content are created. This would result in faster performance for the clients and better control of the bandwidth utilization. However, keeping multiple copies of the same content will reduce the amount of storage available for caching unique content.

To configure OMD to cache more than one copy of the content that is served by the Delivery Service, from the Initial Dispersion drop-down list choose the number of copies of the content that you want to cache. These copies would be stored on different cache servers in the cache group. For example, if you choose 3, a copy of the content is stored on 3 different cache servers.

- **IPv6 Routing Enabled?:** This setting determines whether the Traffic Router will respond to IPv6 DNS (AAAA record) requests for the hostnames of the Delivery Service. The default setting is No, which means the Traffic Router will only respond to IPv4 DNS (A record) requests for the hostnames of the Delivery Service. If you change this setting to Yes, the Traffic Router will respond to both IPv6 DNS (AAAA record) requests and IPv4 DNS (A record) requests for the hostnames of the Delivery Service.
- **Anonymous Blocking Enabled?:** This setting determines whether this Delivery Service will use the Anonymous Blocking feature to determine whether a client request is allowed. If this value is set to Yes, the Anonymous Blocking profile settings that are configured for the Traffic Router that services this Delivery Service will be used to help determine whether a client request is allowed. For information on how Anonymous Blocking works and how it interacts with the other client blocking methods, see [Understanding CDN Client Blocking Options, on page 23](#). For information on how to configure the Anonymous Blocking settings of the Traffic Router profile, see [Configure Anonymous Blocking, on page 169](#).



Note You cannot enable Anonymous Blocking for DNS Delivery Services.

- **Range Request Handling:** Some of the content that the CDN may deliver can be very large, however the client may only need a portion of that file because they are only watching a portion of the content.

To enable a client to request only a portion of a file, browsers and video players support a Range Request header, which denotes what part (bytes) of the content the client is requesting. For example, a client could make a request that they only want bytes 0 through 199 of the content (the first 200 bytes) or they may request that they want bytes from a range later in the file such as bytes 10,000 and 11,000. The Range Request Handling setting determines how the cache server handles the content from these requests. The options for this setting are:

- **0- Don't Cache:** When you choose this option, the cache servers will not cache the content for any requests that contain a range request (essentially any partial object request). The cache servers pass

through the request in order to delivery the content to the client, but they do not cache the content that they receive from the Mid cache group or Origin Servers. This is the default setting.

- **1 - Use background_fetch plugin:** When you choose this option, if a client makes a range request, for example they request only bytes 200-300 of an object, the cache server will cache the entire object but will return to the client only the requested range. The cache server delivers the range to the client as soon as the range is downloaded; it does not have to wait for the entire object to be cached before it can deliver the requested range to the client. However, the cache server will save the entire object in cache, in the background, in case another client wants either a portion of this object or the whole object.
- **2 - Use cache_range_requests plugin:** When you choose this option, the cache server will store requests for the same object but different byte ranges as separate unique objects in cache. For example, if a client asks for a byte range of 0-10 of a certain object, the cache server will cache those bytes as one object. If another client asks for a different byte range for the same object, for example bytes 5-10, the cache server will store this request as a completely separate unique object in cache, even though the byte range overlaps with a previous request.



Note If you have clients that make range requests and you can control these requests so that they are aligned (the clients are requesting the same ranges for the same object), choosing option “2 - Use cache_range_requests plugin” would be beneficial. However, if you have many clients that make range requests but those range requests are not aligned, option 2 is not efficient because having a large number of range requests that are not aligned can result in a large number of objects with overlapping content stored in cache

- **Delivery Service DNS TTL:** This setting determines the TTL, in seconds, that the Traffic Router will set on the A record responses to the clients for the hostnames of this Delivery Service. When the IPv6 Routing Enabled setting is set to Yes, this TTL also applies to the AAAA record responses. The default is 3600 seconds.
- **Geo Miss Default Latitude and Geo Miss Default Longitude:** What these settings determine, depends on the value of the Geo Limit setting:
 - **If the Geo Limit field is set to None:** When Geo Limit is set to None, the Traffic Router will allow all client requests through and will use the CZF file to determine which cache group to direct the client to. If the check of the CZF file does not return a Cache to use for the request, Proximity routing is checked next, if it is enabled in the profile of the Traffic Router. (See [Proximity Routing, on page 159.](#))

If Proximity routing is not enabled, or if Proximity routing does not return a Cache to use for the request, the Traffic Router will use geolocation based routing to determine a cache group the client request should use and then chooses a cache from that group. If no geolocation is available for the requesting address, the Geo Miss Default Latitude and Geo Miss Default Longitude settings of the Delivery Service are used to determine the closest cache group to use for the client request and then chooses a cache from that cache group. (See [Geolocation Based Routing, on page 161.](#))

 - **If the Geo Limit field is set to CZF Only:** When the Geo Limit is set to CZF file Only, the Geo Miss Default Latitude and Geo Miss Default Longitude settings are not used. This is because with this setting, only clients whose IP addresses match an entry in the CZF file are allowed through. If

the client matches an entry in the CZF file, that entry will determine which cache group the client is directed to.

- **If the Geo Limit field is set to CZF + CountryCode(s):** When Geo Limit is set to CZF file + CountryCode(s), if the client IP address does not match an entry in the CZF file, then the Traffic Router tries to determine the location of the client IP address using a geolocation database. If the Traffic Router cannot find the location of the client using the geolocation database, it will then use the Geo Miss Default Latitude and Geo Miss Default Longitude fields to determine what country the request is from.

If the country location of the client matches a country listed in the Geo Limit Country Codes field, the Traffic Router will then check the client request against the Anonymous Blocking configuration to determine whether the request is allowed. If the client request is allowed, the Traffic Router uses Proximity routing, if it is enabled, to determine which cache group to use. (See [Proximity Routing](#), on page 159.)

If Proximity routing is not enabled, or if Proximity routing does not return a Cache to use for the request, the Traffic Router will use the Geo Miss Default Latitude and Geo Miss Default Longitude settings of the Delivery Service to determine the closest cache group to use for the client request. (See [Geolocation Based Routing](#), on page 161.)

The following is an example of the general advanced settings:

The screenshot shows a configuration window titled 'Additional Info' with a 'General' tab. The settings are as follows:

General:	
Query String Handling	0 - use qstring in cache k
Initial Dispersion	1
Regex Remap Expression	
IPv6 Routing Enabled?	No
Geo Limit	None
Anonymous Blocking Enabled?	No
Geo Limit Country Codes	-- Select Countries --
Range Request Handling	0 - Don't cache
Bypass FQDN	
Delivery Service DNS TTL	
Maximum Bits per Second allowed globally (4T or 500M are valid entries)	
Geo Miss Default Latitude	-90.0 to 90.0
Maximum Transaction allowed globally	
Geo Miss Default Longitude	-180.0 to 180.0
Traffic Router Log Request Headers	

Advanced Settings: Content Preposition

- **DSCP Edge Tag:** All of the packets that are sent from the Edge cache to the client can be marked with a DSCP value. The DSCP Edge Tag setting enables you to configure what that DSCP value will be. The DSCP Edge Tag can be used by routers and other network devices to give the traffic from this Delivery Service different priority.

- **DSCP Mid Tag:** All of the packets that are sent from the Mid cache to the Edge cache can be marked with a DSCP value. The DSCP Mid Tag setting enables you to configure what that DSCP value will be. This tag can be used by routers and other network devices to give the traffic from this Delivery Service different priority.
- **Use Content Prepositioning:** Typically content is not acquired and stored on the cache until a client requests the content. However, in some situations, for example if you have content that you know is going to be very popular, you may want the cache to acquire and store the content before it is requested. This is what the Content Prepositioning feature enables you to do.

To configure the Delivery Service to support content prepositioning, choose **Yes** for the Use Content Prepositioning setting. In addition, you must also configure the Ingest Manifest File URL field and the Preload On field for content prepositioning to occur. These fields are discussed below. The default for this setting is No.



Note The following fields are only available to edit if Use Content Prepositioning is set to **Yes**.

- **Ingest Manifest File URL:** When you configure prepositioning, you use an Ingest Manifest file to define what content to put on the selected caches ahead of time. In the Ingest Manifest File URL field, enter the web address (URL) for the location of the Ingest Manifest file. For example, <http://store.company.example/config/cdnmanifest.json>.



Note For an example of an Ingest Manifest file and for a description of its syntax, see [Example Ingest Manifest File, on page 261](#).

- **HTTP Proxy for Ingest Manifest File URL:** If the cache servers need to go through a proxy server to retrieve the Ingest Manifest file configured in the Ingest Manifest File URL field, enter the URL for the proxy server in this field.
- **Preload on:** This field enables you to choose the type of cache servers that you would like preposition the content on. Your options are Edge Only, Mid Only, or Edge + Mid.



Note The following two fields are optional fields that you can configure for Content Prepositioning. These features require additional configuration outside of OMD Director in Traffic Ops.

- **Use Dedicated Volume:** By default cache servers have two volumes:
 - One volume for live content for Live Delivery Services, which is stored in RAM
 - One volume that is typically all of the disks in the cache server, which is used for VOD Delivery Services

Additionally, a third volume can be created that is dedicated to storing prepositioned content for a Delivery Service.

If a third volume has been created on the cache server, the Use Dedicated Volume field points to that volume, which tells the cache server to use that volume for prepositioned content. With this field you also configure which Origin Servers in the Delivery Service you want to preposition content from. The following is the syntax for this field:

hostname=server.company.example volume=X

where *server.company.example* is the URL of the Origin Server whose content you want to preposition on this volume and *X* is the volume number that was created for the dedicated volume. This number must match the Preposition Volume number that was created in Traffic Ops. For example **hostname=s1.company.example volume=3**.

- **Cache Configuration:** By default when content is prepositioned on the cache server, the content stays on the cache until the cache server needs space to hold new content. When the cache fills up and clients make new requests to the cache, the cache automatically evicts the oldest content, which is not always the behavior that you want. When you preposition content in the cache, you might want to ensure that the prepositioned content is always available and that it is never evicted from the cache. The pinning feature enables you to accomplish this.

If the cache server has been configured to support pinning in Traffic Ops, you use the Cache Configuration field to configure which Origin Servers you will pin content for and for how long the content will be pinned. The following is the syntax for this field:

dest_host =server.company.example pin-in-cache=time

where *server.company.example* is the URL of the Origin Server whose content you want to pin and *time* is the length of time to pin this content. For example **dest_host=os1.compay.example pin-in-cache=100d**. In this example 100d means to pin the content for 100 days. For pin-in-cache you can reference “d” for day, “h” for hours, and “m” for minute. If the hostname in the Cache Configuration field matches the hostname in the Use Dedicated Volume field, content for this Origin Server will be pinned to the volume configured in the Use Dedicated Volume field.

- **Session Tracking:** Session tracking assigns all CDN viewers tracking values within an HTTP cookie. Those tracking values are written into the OMD Transaction Logs for further analysis.

To enable session tracking, from the Session Tracking drop-down list choose **Yes**.



Note Session tracking is currently only supported on Edge caches and Mid caches. Traffic Router does not support session tracking.

Session tracking occurs before header rewrite. When any header rewrite rules are defined to rewrite the “Cookie”/ “Set-Cookie” field, carefully define the rule so that it does not interfere with session tracking or the Session Tracking Security Enablement feature.

- **Session Tracking Query Key List:** Session tracking can use either a randomly generated ID or can accept customer parameters from a URL query string. To use customer parameters, in the Session Tracking Query Key List field, enter a comma separated list of URL Query Parameters to capture in the cookie.
- **Session Tracking Security Enablement:** To enable the Session Tracking Security Enablement feature, in the Session Tracking Query Key List field enter **key=<keyvalue>**, where <keyvalue> is the secret key to use to create the cookie signature. If URL Query Parameters also need to be listed in the Session

Tracking Query Key List field for session tracking, add the **key=<keyvalue>** at the end separated by a comma. For example: **sessionId, key=mykey**.



Note To use the Session Tracking Security Enablement feature, Session Tracking must be set to “Yes” and you must have the URL Signing feature configured. For more information on configuring the URL Signing feature, see [URL Signing, on page 127](#).

For more information on Session Tracking Security Enablement, see [Session Tracking Security Enablement, on page 132](#).

The following is an example of the Content Preposition settings:

Content Preposition:	
DSCP Edge Tag	0 - Best Effort
DSCP Mid Tag	--
Use Content Prepositioning	Yes
Ingest Manifest File URL *	http://store.cdn.company
HTTP Proxy for Ingest Manifest File URL	http://proxy.company.exz
Preload on	edge only
Use Dedicated Volume	hostname=s1.company.t
Cache Configuration	dest_host=os1.company
Session Tracking	Yes
Session Tracking Query Key List	sid, uid

Advanced Settings: Header Rewrite Configurations

In addition to remapping a client request to an origin server, there may be a need to modify information in either the client request or the response header. The Header Rewrite feature enables you to define rules to modify the headers for both requests and responses. The rules that you define will determine which requests or responses will be modified, based on matching conditions, and how the header will be modified, based on the action (operator) you define.



Note These header rewrite rules are configured on the Delivery Service and will apply to all of the Edge caches, Mid caches, or Traffic Routers that service that Delivery Service depending on the type of header rewrite rules you define.

- **Edge Header Rewrite Rules:** The Edge cache will modify either the request or response header, based on the rule you define. Enter the rewrite rules in the text box. To enter more than one rewrite rule, separate them by commas. Refer to [Header Rewrite Rules Syntax, on page 267](#) for information on the syntax for this rule.
- **Mid Header Rewrite Rules:** The Mid cache will modify either the request or response header, based on the rule you define. To enter more than one rewrite rule, separate them by commas. Refer to [Header Rewrite Rules Syntax, on page 267](#) for information on the syntax for this rule.



Note Mid Header Rewrites rules are applied based on the Origin FQDN. If multiple Delivery Services share the same origin, mid header rewrites may conflict with each other depending on the contents.

- **Cache URL Expression:** The Cache URL Expression field enables you to change the cache key that is used for caching a request. This enables different URLs that provide the same content to use the same cached object. The format for the entries in this field is:

<pattern> <replacement>

- *<pattern>* is a regular expression that will be applied against the incoming request URL to determine which URLs to match.
- *<replacement>* is the cache key to use for the incoming request URLs that match *<pattern>*.

Example:

`http://s[123].example.com/(.*) http://s.example.com.TSINTERNAL/$1`

With this example, the s1.example.com, s2.example.com, and s3.example.com domains will effectively share the same cache objects.



Note To enter more than one URL cache expression, enter each expression on a separate line.

- **Traffic Router Additional Response Headers:** The Traffic Router will add additional HTTP headers in the response to the client, based on the information you enter. In the Name field, enter the name of an additional header that you want the Traffic Routers to add and in the Value field, enter the value to assign to the header. You can enter multiple headers by clicking the + button to add a new Name/Value pair.



Note This settings is only supported on HTTP-based Delivery Services. You will not see it if you are configuring the Advanced Settings of a DNS-based Delivery Service.



Note Its preferred to configure any CORS related headers using the CORS section, instead of placing them here.

The following is an example of the Header Rewrite Configurations settings:

Header Rewrite Configurations :

Edge Header Rewrite Rules

`cond %{READ_RESPONSE_HDR_HOOK}
rm-header Set-Cookie`

Mid Header Rewrite Rules

Cache URL expression

Traffic Router Additional Response Headers

Cache-C

no-cache

-

+

Advanced Settings: URL Signing

URL signing cannot be configured through the CDN Wizard so this section will be grayed out. After you have completed the CDN Wizard, you can edit the Delivery Service to configure URL signing. For more information, see [Edit an Existing DNS or HTTP Delivery Service, on page 124](#).

Advanced Settings: Traffic Router CORS Settings for HTTP(s) Content Routing

CORS enables browsers to use a predefined set of headers and methods to describe the set of origins that are permitted to access resources using a web browser. “Origin” as it pertains to CORS and the CORS configuration refers to the scope of authority or privilege used by user agents (e.g. web browser) as specified in an “Origin:” header in the requests. It does not refer to the Origin Server of the CDN.

Configure the following settings to enable and configure the CORS policies that need to be enforced in your CDN. (Media Streamer supports both simple and non-simple (preflight) CORS requests.)












Note The CORS settings are only supported on HTTP-based Delivery Services. You will not see them if you are configuring the Advanced Settings of a DNS-based Delivery Service.

- **Allowed Origins:** Enter a URI that may access the resources of the Delivery Service. For example, `cdn.example1.com` or `https://example2.com:8080`. The Traffic Router compares the value in the Origin header of the client CORS request against this list. If there is a match, the request is allowed. To add additional URIs, click the **Add** (plus) icon. To delete a URI from the list, click the **Delete** (minus) icon.
- **Allow Credentials:** If you want the Traffic Router to allow responses to include headers that have the Access-Control-Allow-Credentials header set to true, choose **Yes**.
- **Exposed Headers:** Enter the headers that the client is allowed to access. To add additional headers, click the **Add** (plus) icon. To delete a header from the list, click the **Delete** (minus) icon.

- **Preflight Allowed Methods:** Check the check boxes for the methods that the Traffic Router should allow when a client is accessing the resources of the Delivery Service. This check occurs during a preflight check.
- **Preflight Allowed Headers:** Enter the general headers that the client can use when making the actual request. This check occurs during a preflight request. To add additional headers, click the **Add** (plus) icon. To delete a header from the list, click the **Delete** (minus) icon.
- **Preflight Max Age:** Enter the length of time that a client is allowed to cache the results of a preflight request.

Traffic Router CORS Settings for HTTP(s) Content Routing :

Allowed Origins	<input type="text" value="https://example1.company"/>	 	Preflight Allowed Methods	<input checked="" type="checkbox"/> OPTIONS <input checked="" type="checkbox"/> PUT
				<input type="checkbox"/> DELETE <input type="checkbox"/> TRACE
				<input type="checkbox"/> CONNECT
Allow Credentials	<input type="text" value="Yes"/>		Preflight Allowed Headers	<input type="text" value="X-Requested-With"/>
				<input type="text" value="Client-Security-Token"/>
Exposed Headers	<input type="text" value="Content-Length"/>	 		 
	<input type="text" value="Accept-Encoding"/>	 		
			Preflight Max Age	<input type="text" value="300"/> <input type="text" value="Sec"/>



When you are finished editing the Advanced Settings, click **Save**.

Step 5: Review

The next step in the CDN Wizard is to review the configuration of the CDN elements. If you need to make any changes, you can make these changes from the Review page by using Edit button or Delete button. You can also make changes to these elements after you complete the CDN Wizard by going to Provisioning > CDN Overview or Provisioning > Edit CDN.

The following is an example of the Review page:

Step 5 - Review

Instructions

On this page you can review the elements of your CDN. The edit option allows you to make minor changes to the CDN elements. Please go back to respective pages for setting up additional servers, adding cache groups, creating delivery services and changing assignments.

Delivery Services

10

Delivery Service Name	Routing Type	Protocol	Active	Modify
vod-wc	http	http	UP	

Showing 1 to 1 of 1 entries

Previous

1

Next

Mid Cache Groups

10

Group Name	Number of Servers	Longitude	Latitude	Modify
mid	0	47	-34	

Showing 1 to 1 of 1 entries

Previous

1

Next

Edge Cache Groups

10

Group Name	Number of Servers	Longitude	Latitude	Parent Cache Group	Modify
edge	2	12	32	mid	

When you are finished reviewing the configuration of the CDN elements, click **Next** or click **6 Accept**.

Step 6: Accept

The final step of the CDN Wizard is to accept the configuration. To accept the configuration, click **Finish**. This will execute the configuration of all of the elements that you created during the CDN Wizard. The following is an example:

CDN Wizard

1

2

3

4

5

6

Create CDN

Prepare Servers

Assign Servers

Create Delivery Services

Review

Accept

Step 6 - Accept

Instructions

Click Finish for Initializing CDN and Provisioning CDN elements. The process will take few minutes. During this time please do not click back or refresh your browser. Upon completion you will be taken to CDN Overview page.

CDN Name : sevt-cdn

CDN Status :

Previous

Finish

**Note**

If you receive the error message “Please select profile for all Assigned Servers”, return to Step 5 of the CDN Wizard and in the “Assigned Servers” section, choose the correct profile for each Server based on its type.

After you click Finish, a window appears that shows the progress of the CDN being created, including the status of each action involved. When the CDN has been successfully created, you will see a green check mark for every action in the Status column and the CDN Status will show Success.

Step 6 - Accept

Instructions

Click Finish for Initializing CDN and Provisioning CDN elements. The process will take few minutes. During this time please do not click back or refresh your browser. Upon completion you will be taken to CDN Overview page.

CDN Name : sevt-cdn ✓			
CDN Status : Success			
Action Name	succeeded	Failed	Status
Get Salt Minion Details	edge-cache1, traffic-router1, traffic-monitor1, mid-cache1, edge-cache2		✓
Create CDN	sevt-cdn		✓
Create Monitor Servers	traffic-monitor1		✓
Create Router Servers	traffic-router1		✓
Create Cache Groups	mid1, edge1		✓
Create Cache Servers	edge-cache1, edge-cache2, mid-cache1		✓
Update allow IPs in edge profiles	["sevt-cdn3", "EDGE1_CDN_520"]		✓
Create Delivery Services	vod		✓
Initial Activate Configuration on Router	sevt-cdn3		✓
Salt Server Init Monitors	traffic-monitor1		✓
Salt Server Init Routers	traffic-router1		✓
Execute Operational Readiness Script on Caches	edge-cache1, edge-cache2, mid-cache1		✓
Final Activate Configuration on Router	sevt-cdn3		✓

Upon successful creation of the CDN, you will be taken to the CDN Overview page. From the CDN Overview page you can validate the configuration of the Delivery Services, Mid cache groups, Edge cache groups, and servers that you have provisioned in the CDN. From this page you can also delete these items. For more information on validating this information, see [Validating the CDN Configuration, on page 53](#). If there is any problem creating the CDN, an Alert window appears showing you the cause of the failure.

**Note**

After the CDN Wizard has successfully completed, the servers that were provisioned by the wizard will be in an Offline state.

Validating the CDN Configuration

To validate the configuration of the CDN, choose **Provisioning > CDN Overview**. From the CDN Overview page you can validate the configuration for the following:




- Delivery Services
- Mid cache groups
- Edge cache groups
- Servers

View Delivery Services

To view a list of Delivery Services for a CDN, from the Provisioning > CDN Overview window, choose the CDN from the **CDN Name** drop-down list. The top section of the Provisioning > CDN Overview page lists the Delivery Services that are configured in the CDN, including the following information:

- **Delivery Service Name:** Name of the Delivery Service
- **Routing Type:** DNS, HTTP, STEERING, CLIENT_STEERING
- **Protocol:** HTTP, HTTPS, HTTP and HTTPS, or HTTP to HTTPS
- **Active:** Whether the Delivery Service is active or not
- **Status:** Status of the delivery service

You can sort the list by any of these columns by clicking the arrow icons in the column header. From this section you can also view additional details about the Delivery Service and delete a Delivery Service. The following is an example:





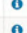
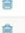






Delivery Services						
<input type="text"/>	<input type="text"/>					
10						
Delivery Service Name	Routing Type	Protocol	Active	Status	Action	
vod	http	http	UP	success	 	
https-dscp	http	https	UP	success	 	

If the word “updating” appears instead of the Delete icon (trashcan) in the Action column, this indicates that the Delivery Service has been modified and is in the process of being updated. To see the status of the update, click the **Updating** link.

Delivery Services						
<input type="text"/>	<input type="text"/>					
10						
Delivery Service Name	Routing Type	Protocol	Active	Status	Action	
vod	http	http	UP	updating	 updating	
https-dscp	http	https	UP	success	 	

To view the details about a Delivery Service, follow these steps:

1. In the Delivery Service section, click the **Information** icon (i) for the Delivery Service that you want to view.

Delivery Services						
<input type="text"/>	<input type="text"/>					
10						
Delivery Service Name	Routing Type	Protocol	Active	Status	Action	
sp1-ds1	http	http	DOWN	success	 	
sp1-ds2	http	http	DOWN	success	 	
sp1_steering	STEERING	http	DOWN	success	 	
Steering	STEERING	http	UP	success	 	
sp1_steering	STEERING	http	DOWN	success	 	
vod2	http	http	UP	success	 	

2. In the window that appears you can view the settings for the Delivery Service.

Delivery Service - omd-ds1

Delivery Service Name

omd-ds1

Display Name

omd-ds1

Long Description

Use Multi Site Origin Feature

No

Delivery Service profile

Routing Type

HTTP

Content Type

VOD

Content Scope

National

Edge cache retrieval from origin

False

Protocol

HTTP

Origin Server Base URL

http://sp1-os1-video.companyx.com

Customer

Active

Yes

Assigned Cache Servers

omd-edge1

Assigned Device Groups

Delivery Service Server Assignments:

Assign Cache Servers

Assign Cache Servers IP

omd-edge1

10.63.231.34 - Primary

Domains for this Delivery Service:

Sub Domain *

Sample Client URL

ds1

https://tr.ds1.spcdn.companyx.com

3. To see the advanced settings for the Delivery Service, click **Show Advanced Settings**.

4. To switch back to the initial settings page, click **Hide Advanced Settings**.

5. When you are done viewing the settings, click **Close**.



Note For information on how to delete a delivery service, see [Delete a Delivery Service, on page 126](#).

View Mid Cache Groups

The Mid Cache Groups section of the Provisioning > CDN Overview page lists the Mid cache groups that are configured in the CDN, including the following information:













Note Make sure the correct CDN is selected from the CDN Name drop-down list.

- Group name
- Latitude and longitude of the Mid cache group
- Status of the group

You can sort the list by any of these columns by clicking the arrow icons in the column header. From this section you can also view additional details about the Mid cache groups or delete a Mid cache group.

Mid Cache Groups

10

Group Name	Latitude	Longitude	Status	Action
MID-SJC	47	-122	success	 
MID-NTN	32.33	34.85	success	 
MID-SG	1.28	103.85	success	 
MID-SJ	37.23	-121.86	success	 
MID-US-VA	36.9	-76.2	success	 

Showing 1 to 5 of 5 entries










Previous1Next



Note If the word “updating” appears instead of the Delete icon (trashcan), this indicates that the Mid cache group has been modified and is in the process of being updated. To see the status of the update, click the **Updating** link.

Mid Cache Groups

10

Group Name	Latitude	Longitude	Status	Action
MID-SJC	47	-122	success	 
MID-NTN	32.33	34.85	success	 
MID-SG	1.28	103.85	success	 
MID-SJ	37.23	-121.86	updating	 updating
MID-US-VA	36.9	-76.2	success	 

Showing 1 to 5 of 5 entries

Previous1Next

To view the details of a Mid cache group, follow these steps:

1. In the **Mid Cache Groups** section, click the Information icon (i) for the Mid cache group that you want to view.

Mid Cache Groups

10

Group Name	Latitude	Longitude	Status	Action
MID-SJC	47	-122	success	<div><div></div><div></div></div>
MID-NTN	32.33	34.85	success	<div><div></div><div></div></div>
MID-SG	1.28	103.85	success	<div><div></div><div></div></div>
MID-SJ	37.23	-121.86	success	<div><div></div><div></div></div>
MID-US-VA	36.9	-76.2	success	<div><div></div><div></div></div>

Showing 1 to 5 of 5 entries

Previous

1

Next

- In the **Mid Cache Group** window, you can view the settings for the Mid cache group.

Mid Cache Group - MID-SJ

Group Name	MID-SJ
Short Name	MID-SJ
Type	MID
Latitude	37.279518
Longitude	-121.867905
Assigned Servers	omd-mid-01.spcdn.company.example
Parent Cache Group	NO_PARENT
Secondary Parent Cache Group	NO_PARENT

Close

- When you are done viewing the details of the Mid cache group, click **Close**.



Note For information on how to delete a Mid cache group, see [Delete a Mid Cache Group, on page 88](#).

View Edge Cache Groups

The Edge Cache Groups section of the Provisioning > CDN Overview page lists the Edge cache groups that are configured in the CDN, including the following information:










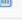


Note Make sure the correct CDN is selected from the **CDN Name** drop-down list.

- Group name
- Latitude and longitude of the Edge cache group
- Parent cache group




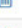

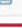

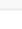
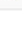
- Status of the group

You can sort the list by any of these columns by clicking the arrow icons in the column header. From this section you can also view additional details of an Edge cache group or delete an Edge cache group.

Edge Cache Groups						
<input type="text"/>		10				
Group Name	Latitude	Longitude	Parent Cache Group	Status	Action	
EDGE-AMS	52.379	4.89	MID-NTN	success	 	
EDGE-JRS	31.9	35.2	MID-NTN	success	 	
EDGE-KJK	50.82	3.26	MID-NTN	success	 	
EDGE-NTN	32.33	34.85	MID-NTN	success	 	
EDGE-SG	1.28	103.85	MID-SG	success	 	






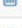






Note If the word “updating” appears instead of the Delete icon (trashcan), this indicates that the Edge cache group has been modified and is in the process of being updated. To see the status of the update, click the updating link.

Edge Cache Groups						
<input type="text"/>		10				
Group Name	Latitude	Longitude	Parent Cache Group	Status	Action	
EDGE-AMS	52.379	4.89	MID-NTN	success	 	
EDGE-JRS	31.9	35.2	MID-NTN	success	 	
EDGE-KJK	50.82	3.26	MID-NTN	success	 	
EDGE-NTN	32.33	34.85	MID-NTN	updating	 updating	
EDGE-SG	1.28	103.85	MID-SG	success	 	

To view the additional details about an Edge cache group, follow these steps:

1. In the **Edge Cache Groups** section, click the Information icon (i) for the Edge cache group that you want to view.

Edge Cache Groups						
<input type="text"/>		10				
Group Name	Latitude	Longitude	Parent Cache Group	Status	Action	
EDGE-AMS	52.379	4.89	MID-NTN	success	 	
EDGE-JRS	31.9	35.2	MID-NTN	success	 	
EDGE-KJK	50.82	3.26	MID-NTN	success	 	
EDGE-NTN	32.33	34.85	MID-NTN	success	 	
EDGE-SG	1.28	103.85	MID-SG	success	 	

2. In the **Edge Cache Group** window, you can view the settings for the Edge cache group.

Edge Cache Group - EDGE_JRS	
Group Name	EDGE-JRS
Short Name	EDGE-JRS
Type	EDGE
Latitude	31.9
Longitude	35.2
Assigned Servers	omd-edge-01-sj.spcdn.company.example
Parent Cache Group	MID-NTN
Secondary Parent Cache Group	NO_PARENT

Close

3. When you are finished viewing the information for the Edge cache group, click **Close**.


Note

For information on how to delete a Mid cache group, see [Delete an Edge Cache Group, on page 89](#).

View Servers









The **Servers** section of the Provisioning > CDN Overview page lists the Edge cache, Mid cache, Origin, Traffic Router, and Traffic Monitor servers that are configured in the CDN, including the following information about these servers:


Note

Make sure the correct CDN is selected from the **CDN Name** drop-down list.








- **Server name:** FQDN of the server
- **IP address:** IP address of the server
- **Type:** Edge, mid, monitor, router, or origin
- **Cache Group:** Name of the cache group to which the server belongs. Servers that show a cache group of OrphanGroup are not currently assigned to a cache group.
- **Profile:** The profile that has been assigned to the server.
- **Status:** Status of the server.

You can sort the list by any of these columns by clicking the arrow icons in the column header. From this section you can also view additional details about the server or delete a server.

Servers							
<div> <input type="text"/> 10 </div>							
Server Name	IP Address	Type	Cache Group	Profile	Status	Action	
omd-edge-01-sj.spcdn.company.example	10.63.231.34	edge	edge2	EDGE-SEA	success	 	
omd-mid-01.spcdn.company.example	10.63.231.35	mid	mid2	MID-DEN	success	 	
omd-monitor.spcdn.company.example	10.63.231.32	monitor	CtrlPlaneGroup	RASCAL_CDN1_sevt-cdn	success	 	
omd-tr.spcdn.company.example	172.24.22.88	router	CtrlPlaneGroup	CCR_CDN_sevt-cdn	success	 	











Note If the word “updating” appears instead of the Delete icon (trashcan), this indicates that the Server has been modified and is in the process of being updated. To see the status of the update, click the **Updating** link.

Servers							
<div> <input type="text"/> 10 </div>							
Server Name	IP Address	Type	Cache Group	Profile	Status	Action	
omd-edge-01-sj.spcdn.company.example	10.63.231.34	edge	edge2	EDGE-SEA	success	 	
omd-mid-01.spcdn.company.example	10.63.231.35	mid	mid2	MID-DEN	updating	 updating	
omd-monitor.spcdn.company.example	10.63.231.32	monitor	CtrlPlaneGroup	RASCAL_CDN1_sevt-cdn	success	 	
omd-tr.spcdn.company.example	172.24.22.88	router	CtrlPlaneGroup	CCR_CDN_sevt-cdn	success	 	

To view additional settings for a server, follow these steps:

1. In the **Servers** section, click the Information icon (i) for the server that you want to view additional details for.

Servers							
<div> <input type="text"/> 10 </div>							
Server Name	IP Address	Type	Cache Group	Profile	Status	Action	
omd-edge-01-sj.spcdn.company.example	10.63.231.34	edge	edge2	EDGE-SEA	success	 	
omd-mid-01.spcdn.company.example	10.63.231.35	mid	mid2	MID-DEN	success	 	
omd-monitor.spcdn.company.example	10.63.231.32	monitor	CtrlPlaneGroup	RASCAL_CDN1_sevt-cdn	success	 	
omd-tr.spcdn.company.example	172.24.22.88	router	CtrlPlaneGroup	CCR_CDN_sevt-cdn	success	 	

2. In the window that appears, you can view the settings for the Server.

Servers - edge-cache1

Server Name	edge-cache1
Type	edge
Profile	Edge_CDN1_Profile
Assigned Device Groups	
Domain Name	crdc.com
Cache Group	ALL_EdgeCG
Admin Status	ONLINE
IP Address	10.74.25.142
Netmask	255.255.255.128
Gateway	10.74.25.129
Interface MTU	1500
Interface Name	eth0

Close



Note To change the Edge cache group to which a server belongs, choose **Provisioning > Edit CDN** and click the **Cache Group** tab. For detailed information, see [Edit an Existing Cache Group, on page 86](#).

3. When you are done viewing the details of the Server, click **Close**.



Note For information on how to delete a Server, see [Delete a Server from CDN Overview Window, on page 72](#).

Print or Export CDN Overview Page

To print the **CDN Overview** page, click the Print icon in the upper-right corner of the Provisioning > CDN Overview page. To export the information from the CDN Overview page to a JSON file, click the **Export** icon in the upper-right corner of the Provisioning > CDN Overview page.

Provisioning > CDN Overview

CDN Name : CDN1

Delivery Services

10

Delivery Service Name	Routing Type	Protocol	Active	Status	Action
testhttps	http	https	UP	success	
test	http	http	UP	success	
URL_Sign	http	http	DOWN	success	



CHAPTER 4

Manage CDN Servers

After you have provisioned the CDN, you may need to add servers to the CDN or modify existing servers. From the Servers tab of the Provisioning > Edit CDN window, you can edit all of the servers that you have in your environment, you can register existing Edge and Mid cache servers that were provisioned as part of the initial OMD installation and add them to your CDN, and you can add Origin Servers to your CDN. This chapter describes how to perform these configuration tasks.



Note

If you are running OMD Director in an HA configuration, you can only make configuration changes from the Primary OMD Director instance. To perform the steps in this chapter, make sure you are logging into the Primary OMD Director instance. If you are logged into the Backup OMD Director instance, you will see “This OMD Director is currently running in backup mode” in the header. If you are logged into an OMD Director running in Detached mode because of an HA failure, you will see “This OMD Director is currently running in detached mode”.



Note

To delete a server (except for an Origin server) you must go to Provisioning > CDN Overview. You can delete an Origin Server from either Provisioning > CDN Overview or the Provisioning > Edit CDN.

This chapter includes the following topics:

- [Add an Origin Server to the CDN, on page 63](#)
- [Add a New Registered Server to the CDN, on page 65](#)
- [Edit Traffic Monitor or Traffic Router Servers, on page 66](#)
- [Edit an Existing Edge Cache Server, on page 67](#)
- [Edit an Existing Mid Cache Server, on page 70](#)
- [Edit an Existing Origin Server, on page 71](#)
- [Delete a Server from CDN Overview Window, on page 72](#)
- [Delete an Origin Server from Edit CDN Window, on page 73](#)

Add an Origin Server to the CDN

Origin Servers need to be added to the CDN if:

- You have a Delivery Service that is using MSO

- An Origin Server will be used by more than one Delivery Service

Typically the Origin Servers are not automatically registered with the CDN so you will need to manually add them. When you add an Origin Server to the CDN, you will also assign it an Origin Server Profile to use and you will assign it to an Origin Server cache group.

Perform the following steps to add an Origin Server to OMD Director:

Procedure

-
- Step 1** Choose **Provisioning > Edit CDN**. From the CDN Name drop-down list above the tabs, make sure the correct CDN is selected.
- Step 2** Next to the **Server Name** drop-down list, click the **Add (plus)** icon.
- Step 3** In the Add New Server window that appears, configure the following information to add a new server:
- **Name:** Enter the hostname of the Origin Server. Do not enter the domain. You will enter domain in the Domain Name field.
 - **IP Address:** Enter the IP address of the Origin Server.
 - **Netmask:** Enter the subnet mask of the Origin Server.
 - **Gateway:** Enter the IP address of the default gateway for the Origin Server.
 - **IPv6 Address:** Enter the IPv6 address of the Origin Server.
 - **IPv6 Gateway:** Enter the IPv6 address of the default gateway for the Origin Server.
 - **Cache Group:** Choose the Origin cache group to which the server should be assigned. You can assign the Origin cache groups to a Mid Cache's Parent Cache Group and Secondary Parent Cache Group settings to determine the primary and backup Origin Servers for the Mid caches in that group. For information on creating an Origin cache group, see [Create Origin Server Cache Groups, on page 136](#).
 - **Profile:** Choose the Origin Server profile to assign to this server. Only Origin Server profiles will appear in this list. The profile that is assigned to the Origin Server is used for MSO and determines the rank of that server when a Mid cache group does not have a Parent Cache Group configured or if the Parent Cache Group has more than one Origin Server assigned. If this Origin Server is not serving any Delivery Services that use MSO, you can choose any profile.
 - **Advanced Settings:**
 - **TCP Port:** The default port is 80. If the Origin Server listens on a different port, change this value to the port.
 - **MTU:** If an MTU other than 1500 is needed, enter the new value.
 - **Domain Name:** Enter the domain name of the Origin Server. The values entered in the Name field and the Domain Name field will create the FQDN of the Origin Server.
 - **Admin Status:** If the Origin Server is ready to be online, choose **Online**. Otherwise choose Offline and come back to change this setting once the Origin Server is ready to be online.
- Step 4** Click **Add** to save and add the new server.
- Step 5** When the server has been successfully added, the "Processing" status in the title bar will disappear and you will see a green check mark for every action in the Status column. Click **OK** to close the status window.

- Step 6** After you add the Origin Server, you can optionally assign it to a device group. If you want to assign the Origin Server to a device group, in the **Assign Device Groups** field, choose the device groups to which you would like to add the server. When you are finished, click **Save**.

Add a New Registered Server to the CDN

To add a registered server to the CDN, follow these steps:

Procedure

- Step 1** Choose **Provisioning > Edit CDN**.
- Step 2** Make sure the **Servers** tab is selected.
- Step 3** The bottom section of the **Servers** tab will show servers that have not yet been registered to the CDN. The following is an example:

Registered Servers									
Server Name	Fingerprint	Type	Status	Profile	Cache Group		Delete	Add Server to CDN	
omd-edge-03	fc b0:c8:90:ba:3d:1a:d5:39:62:52:20:0e:f6:28:d2	edge	Unregistered	-- Select Profile--	-- Select Cache Group				
omd-mid-03	15:55:bf:a2:71:f2:90:dd:d1:28:4a:90:0d:44:fe:e6	mid	Unregistered	-- Select Profile--	-- Select Cache Group				

- Step 4** Click the arrow icon for the server you want to add.
- Step 5** From the **Profiles** drop-down list, choose a profile for each server. These profiles were created as part of the initial OMD installation. Only the profiles that are appropriate for the server, based on the server type, will appear in the drop-down list.
- Step 6** If you are adding a cache server, from the **Cache Group** drop-down list, choose the cache group to add this server to. Only groups that are appropriate for the cache type will appear in the drop-down list.

Registered Servers									
Server Name	Fingerprint	Type	Status	Profile	Cache Group	Install	Delete	Add Server to CDN	
edge-4. companyx.com	d7:f7:52:5f:b3:3f:7b:84:fd:6a:03:46:8c:35:21:54	edge	Unregistered	EDGE_AT5_62	EdgeCG				
edge-5. companyx.com	af:22:1c:32:cc:b6:54:16:c4:1b:e4:c4:66:68:c8:7c	edge	Unregistered	-- Select Profile--	-- Select Cach				

- Step 7** When the status of the server shows “Provision Success”, the Add icon at the end of the row is enabled. Click the Add icon to finishing adding the server to the CDN.

Registered Servers									
Server Name	Fingerprint	Type	Status	Profile	Cache Group	Install	Delete	Add Server to CDN	
edge-cache4.ordc.com	9b:45:d6:e7:d5:de:b2:5c:11:d2:e4:03:3a:cb:45:59		Provision Success	EDGE_AT5_62	EdgeCG				
edge-cache5.ordc.com	15:55:bf:a2:71:f2:90:dd:d1:28:4a:90:0d:44:fe:e6		Unregistered	-- Select Profile--	-- Select Cache Group				

- Step 8** To confirm that the server was added to the CDN, you can confirm that it appears in the Server Name list.

- Step 9** After you add the server to the CDN, you can optionally assign it to a device group. If you want to assign the server to a device group, in the **Assign Device Groups** field, choose the device groups to which you would like to add the server. When you are finished, click **Save**.

Edit Traffic Monitor or Traffic Router Servers

To edit an existing Traffic Monitor or Traffic Router server, follow these steps:

Procedure

- Step 1** Choose **Provisioning > Edit CDN**. From the CDN Name drop-down list above the tabs, make sure the correct CDN is selected.
- Step 2** From the Servers tab, choose a Traffic Monitor or Traffic Router server from the Server Name list.
- Step 3** Change the desired settings. From the Servers tab you can change the following settings of the Traffic Monitor/Router server:
- **IP Address:** This is the IP address assigned to the Traffic Monitor or Traffic Router server.
 - **Netmask:** This is the netmask for the primary IP address assigned to the server.
 - **Gateway:** This is the gateway for the primary IP address assigned to the server.
 - **IPv6 Address:** This is the optional IPv6 address assigned to the server.
 - **IPv6 Gateway:** This is the optional IPv6 gateway address assigned to the server.
 - **Ethernet Interface:** This is the primary ethernet interface for the server.
 - **MTU:** This is the MTU of the server interface.
 - **Cache Group:** The cache group to which the server belongs. This is a read-only value.
 - **Profile:** The profile that is assigned to this server. This is a read-only value.
 - **Admin Status**

- Step 4** When you have finished making your changes, click **Save**.
- Step 5** When the server has been successfully updated, the “Processing” status in the title bar will disappear and you will see a green check mark for every action in the Status column. Click **OK** to close the status window.

Edit an Existing Edge Cache Server

To edit an exiting server, follow these steps:

Procedure

- Step 1** Choose **Provisioning** > **Edit CDN**. From the CDN Name drop-down list above the tabs, make sure the correct CDN is selected.
- Step 2** From the **Servers** tab, choose the Edge cache server from the Server Name list.

The screenshot shows the 'Provisioning > Edit CDN' window. At the top, there's a 'CDN Name' dropdown set to 'omd-cdn1' with a 'Delete' button. Below this are several tabs: 'Servers', 'Cache Groups', 'Delivery Services', 'Client Routing', 'Profiles', and 'Content Invalidation'. The 'Servers' tab is active. It contains a 'Server Name' dropdown menu that is open, displaying a list of server names: 'omd-edge2-site2', 'omd-edge1', 'omd-edge1-site2', 'omd-edge2', 'omd-edge-01' (highlighted), 'omd-mid1', 'omd-mid2', 'omd-monitor', 'omd-router', 'omd-vault', 'org-server', 'org-server2', and 'edge-cache'. To the right of the dropdown are input fields for 'IP Address' (255.255.255.128), 'Ethernet Interface' (1500), and 'Cache Group' (EDGE_ATS_622).

- Step 3** Change the desired settings. From the **Servers** tab you can change the following settings of the Edge cache server:

- **IP Address:** This is the primary IP address assigned to the primary streaming interface of the server, which is used by default to stream content. If needed, secondary IP addresses can be added to the Edge cache to use to stream content for certain Delivery Services. Secondary IP addresses can be assigned to either the same interface that the primary IP address is assigned to, or to a different interface that is available for streaming content. For information about adding a secondary IP address to use for streaming content, see [Manage Secondary Streaming IPs, on page 68](#).

Note Secondary IP addresses can only be added for Edge cache servers.

- **Netmask:** This is the netmask for the primary IP address assigned to the server.
- **Gateway:** This is the gateway for the primary IP address assigned to the server.
- **Ethernet Interface:** This is the primary interface that the server uses to stream content. This is the interface to which the primary IP address is assigned.
- **MTU:** This is the MTU of the primary streaming interface.

- **Cache Group:** The cache group to which the server belongs. The server can only belong to one cache group.
- **Profile:** The profile that is assigned to this server.
- **Admin Status**

Step 4 When you have finished making your changes, click **Save**.

Step 5 When the server has been successfully updated, the “Processing” status in the title bar will disappear and you will see a green check mark for every action in the Status column. Click **OK** to close the status window.

Manage Secondary Streaming IPs

Edge cache servers can stream content using primary or secondary IP addresses and you can identify which IP address the Edge cache should use based on the Delivery Service to which it is assigned. The secondary IP addresses can be assigned to either the same interface as the primary IP address or a different interface.

To add a secondary IP address to the Edge cache server that can be used to stream content, perform the following steps:

Procedure

- Step 1** Choose **Provisioning > Edit CDN**. From the CDN Name drop-down list above the tabs, make sure the correct CDN is selected.
- Step 2** Click the **Servers** tab and from the Server Name drop-down list, choose the Edge cache for which you want to add a secondary IP address.
- Step 3** Expand the **Manage Secondary Streaming IP** area. It shows the currently configured interfaces and any secondary IP addresses that are configured.

- Step 4** You can assign a secondary IP address to use to stream content to the same interface that the Primary IP address is assigned to or to a different interface that is available for streaming. To assign the secondary streaming IP address to a different interface, click the + icon in the Server Interfaces section to add an interface in OMD Director for this Edge cache.

Note If you are assigning the secondary streaming IP address to the primary streaming interface, go to Step 4.

Step 5 In the Add Secondary Streaming Interface dialog box that appears, configure the following fields and then click **Add** to add the interface:

- **Interface:** From the Interface drop-down list, choose the name of a second interface on the Edge cache server to use for streaming content. If you do not see the interface listed, choose Add New and enter the name of the second interface that is available on the Edge cache server to stream content. What you enter in this field *must* match the name of an interface in Linux on the Edge cache server.

Warning If you enter an interface name that does not exist on the server, the health check for the interface will fail and the Edge cache will be marked as unavailable. This will prevent the Traffic Router from redirecting requests to this Edge cache.

- **MTU:** Choose either 1500 or 9000.

Step 6 To add a secondary streaming IP address and assign it to an interface, click the + icon in the Secondary Streaming IPs section.

Step 7 In the Add Secondary IP pop-up window that appears, enter the following information and then click **Add** to add the IP address:

- **Ethernet Interface:** Choose the interface to which you want to assign the secondary IP address.
- **IP Stack:** Choose whether you want to add only an IPv4 secondary IP address, only an IPv6 secondary IP address, or both an IPv4 and IPv6 secondary IP address.
- **IP Address:** From the IP Address drop-down list, choose the secondary IP address to use for streaming content. If you do not see the IP address listed, choose Add New and enter the secondary IP address to use for streaming content. The IP address you enter *must* match an address already configured on the chosen interface in Linux. Entering an address in this field will *not* assign the address to the underlying interface.

Warning If you enter an IP address that does not match an address already configured on the chosen interface, the health check for the secondary IP address will fail and the Edge cache will be marked as unavailable. This will prevent the Traffic Router from redirecting requests to this Edge cache.

- **Netmask:** Enter the subnet mask to use with the secondary IP address. The subnet mask you enter must match the subnet mask already configured on the chosen interface in Linux.
- **Gateway:** This is a required field but is for informational purposes only. This field is not used by Media Streamer. You can leave this field set to the default value, which is the gateway for the primary IP address assigned to the server.

What to do next

To determine which IP address an Edge cache server will use to stream content, you choose the IP address of the Edge cache server either when you assign the server to a Delivery Service or when you assign the Edge cache to a Device Group. For more information on choosing the IP address when assigning the Edge cache to a Delivery Service, see [Add a New DNS or HTTP Delivery Service, on page 92](#) or [Edit an Existing DNS](#)

or [HTTP Delivery Service, on page 124](#). For more information on choosing the IP address when assigning the Edge cache to a Device Group, see [Create a Device Group, on page 182](#) or [Edit a Device Group, on page 184](#).

Edit an Existing Mid Cache Server

To edit an exiting Mid cache server, follow these steps:

Procedure

- Step 1** Choose **Provisioning** > **Edit CDN**. From the CDN Name drop-down list above the tabs, make sure the correct CDN is selected.
- Step 2** From the **Servers** tab, choose the Mid cache server from the **Server Name** list.

The screenshot shows the 'Provisioning > Edit CDN' interface. At the top, there is a 'CDN Name' dropdown set to 'omd-cdn1' and a 'Delete' button. Below this is a tabbed interface with 'Servers', 'Cache Groups', 'Delivery Services', 'Client Routing', 'Profiles', and 'Content Invalidation'. The 'Servers' tab is active, displaying a table with columns for 'Server Name', 'Status', 'IP Address', 'Ethernet Interface', 'Cache Group', and 'Profile'. A dropdown menu is open for the 'Server Name' column, showing a list of servers: 'Orgserver1', 'omd-edge1', 'omd-edge1-site2', 'omd-edge2', 'omd-mid1' (highlighted in blue), 'omd-mid2', 'omd-monitor', 'omd-router', and 'omd-vault'. To the right of the table, there are input fields for 'IP Address' (255.255.255.0), 'Ethernet Interface' (1500), 'Cache Group' (mid-cache-grp), and 'Profile' (MID_ATS_622).

- Step 3** Change the desired settings. From the **Servers** tab you can change the following settings of the Mid cache server:
- **IP Address:** IP address used to stream content. This IP address must already be configured on the server.
 - **Netmask:** Netmask for the IP address that is used to stream content. This must match the netmask already configured on the server.
 - **Gateway:** Gateway for the IP address that is used to stream content. This must match the gateway already configured on the server.

- **Ethernet Interface:** Name of the interface that is used to stream content.
- **MTU:** The MTU can be either 1500 or 9000.
- **Cache Group:** The cache group to which the server belongs. The server can only belong to one cache group.
- **Profile:** The profile that is assigned to this server.
- **Admin Status**

Step 4 When you have finished making your changes, click **Save**.

Step 5 When the server has been successfully updated, the “Processing” status in the title bar will disappear and you will see a green check mark for every action in the Status column. Click **OK** to close the status window.

Edit an Existing Origin Server

To edit an exiting origin server, follow these steps:

Procedure

- Step 1** Choose **Provisioning** > **Edit CDN**. From the **CDN Name** drop-down list above the tabs, make sure the correct CDN is selected.
- Step 2** From the **Servers** tab, choose the origin server from the **Server Name** list.
- Step 3** Change the desired settings. From the **Servers** tab you can change the following settings of the origin server:
- **IP Address:** IP address used to stream content. This IP address must already be configured on the server.
 - **Netmask:** Netmask for the IP address that is used to stream content. This must match the netmask already configured on the server.
 - **Gateway:** Gateway for the IP address that is used to stream content. This must match the gateway already configured on the server.
 - **IPv6 Address:** IPv6 address used to stream content. This IPv6 address must already be configured on the server.
 - **IPv6 Gateway:** Gateway for the IPv6 address that is used to stream content. This must match the gateway already configured on the server.
 - **Ethernet Interface:** Name of the interface that is used to stream content.
 - **MTU:** The MTU can be either 1500 or 9000.
 - **Cache Group:** The cache group to which the server belongs. The server can only belong to one cache group.
 - **Profile:** The profile that is assigned to this server. The profile that is assigned to the Origin Server is used for MSO and determines the rank of that server when a Mid cache group does not have a Parent Cache Group configured or if the Parent Cache Group has more than one Origin Server assigned. If this Origin Server is not serving any Delivery Services that use MSO, you can choose any profile.

- **TCP Port:** This is the port number that the Origin Server listens on.
- **Domain Name:** This is the domain name of the Origin Server. The values entered in the Name field and the Domain Name field will create the FQDN of the Origin Server.
- **Admin Status**

Step 4 When you have finished making your changes, click **Save**.

Step 5 When the server has been successfully updated, the “Processing” status in the title bar will disappear and you will see a green check mark for every action in the **Status** column. Click **OK** to close the status window.

Delete a Server from CDN Overview Window

To delete a server from the CDN Overview window, follow these steps:



Note

If you are running OMD Director in an HA configuration, you can only make configuration changes from the Primary OMD Director instance. To perform the steps in this section, make sure you are logging into the Primary OMD Director instance.

Procedure

Step 1 Choose **Provisioning > CDN Overview**.

Step 2 In the Servers section, click the **Delete** icon (trashcan) for the server that you want to delete.

Note The CDN system requires one Edge cache server, one Mid cache server, one Traffic Router, and one Traffic Monitor.

Note If the word “updating” appears instead of the Delete icon (trashcan), this indicates that the Server has been modified and is in the process of being updated. To delete the Server, you will need to wait until it has finished updating.

Servers							
<div> <input type="text"/> 10 </div>							
Server Name	IP Address	Type	Cache Group	Profile	Status	Action	
omd-edge-01	10.63.231.34	edge	edge2	EDGE-SEA	success		
omd-mid-01	10.63.231.35	mid	mid2	MID-DEN	success		
omd-monitor	10.63.231.32	monitor	CtrlPlaneGroup	RASCAL_CDN1_sevt-cdn	success		
omd-tr	172.24.22.88	router	CtrlPlaneGroup	CCR_CDN_sevt-cdn	success		

Step 3 In the Deleting Server confirmation window that appears, click **Yes** to confirm that you want to delete the server.

Step 4 When the Server has been successfully deleted, the “Processing” status in the title bar will disappear and you will see a green check mark for every action in the **Status** column. Click **OK** to close the status window.

Delete an Origin Server from Edit CDN Window

To delete an Origin server from the Edit CDN window, follow these steps:

**Note**

If you are running OMD Director in an HA configuration, you can only make configuration changes from the Primary OMD Director instance. To perform the steps in this section, make sure you are logging into the Primary OMD Director instance.

Procedure

- Step 1** Choose **Provisioning > Edit CDN**. From the CDN Name drop-down list, choose the CDN you are editing.
- Step 2** Make sure the **Servers** tab is selected.
- Step 3** From the Server Name drop-down list, choose the Origin Server you want to delete.
- Step 4** Click the **OK** button next to the **Server Name** field.
- Step 5** In the Deleting Server confirmation window that appears, click **Yes** to confirm that you want to delete the server.
- Step 6** When the Server has been successfully deleted, the “Processing” status in the title bar will disappear and you will see a green check mark for every action in the Status column. Click **OK** to close the status window.



CHAPTER 5

Manage Cache Groups

After you have provisioned the CDN, you may need to add cache groups to the CDN, modify existing cache groups, or configure additional features that you can not configure using the CDN Wizard, such as backup Edge cache groups. From the Cache Groups tab of the Provisioning > Edit CDN window you can edit all of the Mid, Edge, and Origin cache groups that you have in your environment and you can create new cache groups.

From the Cache Group tab you can add a new cache group, including Origin server cache groups, and change the following settings of a cache group:

- Short name
- Latitude
- Longitude
- Parent and Secondary Parent Cache Group for Edge cache groups and Mid cache groups
- Configure fallback settings for backup Edge cache groups



Note

To assign cache servers and Origin Servers to a group, you must go to the Provisioning > Edit CDN > Servers tab. To delete a Mid cache group or Edge cache group, you must go to Provisioning > CDN Overview.



Note

If you are running OMD Director in an HA configuration, you can only make configuration changes from the Primary OMD Director instance. To perform the steps in this chapter, make sure you are logging into the Primary OMD Director instance.

This chapter describes how to perform these configuration tasks.

This chapter includes the following topics:

- [Add a Edge Cache Group, on page 76](#)
- [Backup Edge Cache Groups, on page 78](#)
- [Add a Mid Cache Group, on page 84](#)
- [Edit an Existing Cache Group, on page 86](#)
- [Delete a Mid Cache Group, on page 88](#)
- [Delete an Edge Cache Group, on page 89](#)

Add a Edge Cache Group

Edge caches provide edge caching, content streaming, and download to subscriber IP devices. Traffic routers redirect client requests to edge caches based on geolocation, server availability, server load, and server cache content to provide efficient system-wide load balancing. Edge caches are organized into cache groups. You configure each edge cache group with a primary mid-tier parent cache group and optionally a secondary mid-tier parent cache group for failover.

To add a new Edge cache group to the CDN, follow these steps:



Note

For information on Origin cache groups, see [Create Origin Server Cache Groups, on page 136](#).

Procedure

- Step 1** Choose **Provisioning** > > **Edit CDN**. From the **CDN Name** drop-down list above the tabs, make sure the correct CDN is selected.
- Step 2** Click the **Cache Groups** tab.
- Step 3** Next to the **Cache Group Name** drop-down list, click the + icon to create a new Edge cache group.

Provisioning > Edit CDN

CDN Name :

Servers **Cache Groups** Delivery Services Client Routing Profiles Content Invalidation

Cache Group Name

Status success

Short Name *

Latitude *

Longitude *

Type *

Parent Cache Group *

Secondary Parent Cache Group *

- Step 4** In the New Cache Group Window that appears, enter the following information for the new group:
- **Name:** Enter a descriptive name for the group. This is the name that appears in the OMD Director Drop-Down menus. The name should begin with EDGE (in any upper case or lower case combination) so you can easily identify the type of the group.

- **Short Name:** Enter an additional descriptive name. This name is not currently used in OMD Director, but it is a required field. The Short Name can be the same as the Name.
- **The Geo Magnetic Latitude and Geo Magnetic Longitude:** These parameters define the geolocation of the cache group. For cache groups, the geolocation is used by the CDN system to help redirect user clients to the most optimal cache. Enter the latitude and longitude that geolocation should use for this group.

Note For more information on how the Traffic Router determines which cache group should service a request, see the [Client Routing Overview](#) section.

- **Type:** Choose **EDGE_LOC**.

Edge caches provide edge caching, content streaming, and download to subscriber IP devices. Traffic routers redirect client requests to edge caches based on geolocation, server availability, server load, and server cache content to provide efficient system-wide load balancing. Edge caches are organized into cache groups. You configure each edge cache group with a primary mid-tier parent cache group and optionally a secondary mid-tier parent cache group for failover.

- **Parent Cache Group:** Choose the Mid cache group to use as a parent.
- **Secondary Parent Cache Group:** To provide failover in case the Mid caches in the Parent Cache Group are unavailable, you can optionally assign a secondary cache group to use. When the Mid caches in the Parent Cache Group are unavailable, the mid caches in the Secondary Parent Cache Group are used as a backup.

Step 5 Click **Add** to add the group.

The screenshot shows the 'Add New Cache Group' form with the following values:

- Name: EDGE
- Short Name: EDGE
- Geo Magnetic Latitude: -23
- Geo Magnetic Longitude: -46
- Type: EDGE_LOC
- Parent Cache Group: mid
- Secondary Parent Cache Group: NO_PARENT

The 'Add' button is highlighted with a red box.

Note To add a server to a cache group, edit the cache server properties. For more information see [Edit an Existing Origin Server, on page 71](#).

Step 6 When the cache group has been successfully added, the “Processing” status in the title bar will disappear and you will see a green check mark for every action in the Status column. Click **OK** to close the status window.

Note After an Edge Cache group is created, you can configure the Edge cache group to use backup Edge cache groups. Backup edge cache groups are used when there are no caches available in the original Edge cache group selected by the Traffic Router for the client request. For more information on how this feature works and how to configure it, see [Backup Edge Cache Groups, on page 78](#).

Backup Edge Cache Groups

Backup Edge cache groups are used when there are no caches available in the original Edge cache group selected by the Traffic Router for a client request. The Traffic Routers use the CZF file and the Fallback configuration on the Edge cache group to support backup Edge cache groups.

When no Edge caches are available from the original cache group that is matched in the CZF file, the Traffic Router uses the processes described in the following sections, "Process When Geo Limit is Set to CZF Only" and "Process When Geo Limit is Set to CZF + Country Codes or None", to determine which Edge cache group to use, if any, as a backup.



Note [CZF File, on page 82](#) explains how to configure the CZF file to support backup Edge cache groups and provides an example of a CZF.

Process When Geo Limit is Set to CZF Only

The Traffic Router checks the CZF file for a cache group with a subnet match to the client request. If a cache group is found, the Traffic Router then chooses an Edge cache from that group based on cache availability, cache load, and cache content. If there is no available Edge cache for the cache group that was matched, the Traffic Router uses the following process to look for a backup Edge cache group:

- **When the original cache group has fallback cache groups assigned:**

1. The Traffic Router checks the list of fallback cache groups in the order listed, to find an available cache.
2. If there are no Edge caches available from the fallback cache groups, the Traffic Router checks the "Fallback Enable" setting of the original Edge cache group that was matched:
 - If the Fallback Enable setting is set to **No**, no further searches are performed and no cache is returned. The client is redirected to the bypass location listed in the Bypass FQDN field of the Delivery Service. If the Bypass FQDN field is empty, the client will receive a 503 Service Unavailable message.
 - If the Fallback Enable setting is set to **Yes**, the Traffic Router looks for the next closest cache group to the original group, based on the latitude and longitude configured in the CZF file. The Traffic Router then chooses an Edge cache from that group based on cache availability, cache load, and cache content. If no Edge cache is available in the backup Edge cache group, or if there are no latitude and longitude settings in the CZF file for any of the other caches assigned to the Delivery Service used for the client request, the client is redirected to the bypass location listed in the Bypass FQDN field of the Delivery Service. If the Bypass FQDN field is empty, the client will receive a 503 Service Unavailable message.

• **When the original cache group does *not* have fallback cache groups assigned:**

1. The Traffic Router looks for the next closest cache group to the original group, based on the latitude and longitude configured in the CZF file. The Traffic Router then chooses an Edge cache from that group based on cache availability, cache load, and cache content. If no Edge cache is available in the backup Edge cache group, or if there are no latitude and longitude settings in the CZF file for any of the other caches assigned to the Delivery Service used for the client request, the client is redirected to the bypass location listed in the Bypass FQDN field of the Delivery Service. If the Bypass FQDN field is empty, the client will receive a 503 Service Unavailable message.



Note The list of fallback cache groups and the Fallback Enable setting are configured on the Edge Cache group. For more information on configuring these settings, see [Configure Backup Edge Cache Groups, on page 80](#).

Process When Geo Limit is Set to CZF + Country Codes or None

The Traffic Router first checks the CZF file for a cache group with a subnet match to the client request. If a cache group is found, the Traffic Router then chooses a cache from that group based on cache availability, cache load, and cache content. If there is no available cache for the cache group that was matched, the Traffic Router uses the following process to look for a backup Edge cache group:

• **When the original cache group has fallback cache groups assigned:**

1. The Traffic Router checks the list of fallback cache groups in the order listed, to find an available cache.
2. If there are no Edge caches available from the fallback cache groups, the Traffic Router checks the "Fallback Enable" setting of the original Edge cache group that was matched:
 - If the Fallback Enable setting is **No**, no further searches are performed and no Edge cache is returned. The client is redirected to the bypass location listed in the Bypass FQDN field of the Delivery Service. If the Bypass FQDN field is empty, the client will receive a 503 Service Unavailable message.
 - If the Fallback Enable setting is **Yes**, the Traffic Router looks for the next closest Edge cache group to the original group, based on the latitude and longitude configured in the CZF file. It then chooses an Edge cache from that group based on cache availability, cache load, and cache content.
3. If the Fallback Enable setting is **Yes** and no Edge cache is available in the backup Edge cache group selected by the CZF file, or if there are no latitude and longitude settings in the CZF file for any other caches assigned to the Delivery Service being used for the client request, the Traffic Router uses geolocation to find an Edge cache group. The Traffic Router checks the geolocation database to determine the latitude and longitude of the IP address in the client request and compares that latitude and longitude to the latitude and longitude configured for each cache group (in the Cache Group settings) to find the geographically closest cache group. The Traffic Router then chooses an Edge cache from that group based on cache availability, cache load, and cache content.
4. If no Edge cache is available from the Edge cache group based on the geolocation look up, the client is redirected to the bypass location listed in the Bypass FQDN field of the Delivery Service. If the Bypass FQDN field is empty, the client will receive a 503 Service Unavailable message.



Note If coordinates for the client request are not available using geolocation, the Geo Miss Default Latitude and Geo Miss Default Longitude settings of the Delivery Service are used to determine the closest cache group to use. The Traffic Router then chooses a cache to use from that group based on cache availability, cache load, and cache content. If no edge cache is available from that group, the client will receive a 503 Service Unavailable message.

• **When the original cache group does *not* have fallback cache groups assigned:**

1. The Traffic Router looks for the next closest Edge cache group to the original group, based on the latitude and longitude configured in the CZF file. It then chooses an Edge cache from that group based on cache availability, cache load, and cache content.
2. If no Edge cache is available in the backup Edge cache group selected by the CZF file, or if there are no latitude and longitude settings in the CZF file for any other caches assigned to the Delivery Service being used for the client request, the Traffic Router uses geolocation to find an Edge cache group. The Traffic Router checks the geolocation database to determine the latitude and longitude of the IP address in the client request and compares that latitude and longitude to the latitude and longitude configured for each cache group (in the Cache Group settings) to find the geographically closest cache group. The Traffic Router then chooses an Edge cache from that group based on cache availability, cache load, and cache content.
3. If no Edge cache is available from the Edge cache group based on the geolocation look up, the client is redirected to the bypass location listed in the Bypass FQDN field of the Delivery Service. If the Bypass FQDN field is empty, the client will receive a 503 Service Unavailable message.



Note If coordinates for the client request are not available using geolocation, the Geo Miss Default Latitude and Geo Miss Default Longitude settings of the Delivery Service are used to determine the closest cache group to use. The Traffic Router then chooses a cache to use from that group based on cache availability, cache load, and cache content. If no edge cache is available from that group, the client will receive a 503 Service Unavailable message.




Note The list of fallback cache groups and the Fallback Enable setting are configured on the Edge Cache group. For more information on configuring these settings, see [Configure Backup Edge Cache Groups, on page 80](#).

Configure Backup Edge Cache Groups

To configure an Edge cache group to use backup cache groups when none of the caches in the primary Edge cache group are available, perform the following steps:

Procedure

- Step 1** Choose **Provisioning** > > **Edit CDN**. From the CDN Name drop-down list above the tabs, make sure the correct CDN is selected.
- Step 2** Click the **Cache Groups** tab.
- Step 3** From the Cache Group Name drop-down list, choose the Edge cache group for which you want to configure backup Edge cache groups.
- Step 4** From the Fallback Enable drop-down list, choose either **Yes** or **No**, to determine whether the CZF file should be used to find a backup cache group:
- **Yes:** If there are no cache groups entered in the "Fallback Cache Group" field or if there are no caches available from any of the backup groups listed, the Traffic Router will look for a backup cache group using the latitude and longitude configured in the CZF file for the cache groups.
 - **No:** If there are no caches available from any of the backup groups listed, the Traffic Router will *not* check the coordinates in the CZF file so the request will fail. The client is redirected to the bypass location listed in the Bypass FQDN field of the Delivery Service. If the Bypass FQDN field is empty, the client will receive a 503 Service Unavailable message.
- Note** When there are no Fallback cache groups assigned in the **Fallback Cache Group** field, the Fallback Enable setting is always considered to be Yes, regardless of what is set in the **Fallback Enable** field.
- Step 5** From the **Fallback Cache Group** field, choose the Edge cache groups to use for backup. The Traffic Router will check these groups in the order in which they are listed.
- The following is an example of the **Cache Groups** tab:

 Provisioning > Edit CDN

CDN Name :

Servers	Cache Groups	Delivery Services	Client Routing	Profiles	Content Invalidation
<div>Cache Group Name <input type="text" value="EDGE-CG1"/> <input type="button" value="+"/></div> <div>Status success</div> <div>Short Name * <input type="text" value="ECG1"/></div> <div>Latitude * <input type="text" value="60"/></div> <div>Longitude * <input type="text" value="160"/></div> <div>Type * <input type="text" value="EDGE_LOC"/></div> <div>Parent Cache Group * <input type="text" value="MID-CG1"/></div> <div>Secondary Parent Cache Group * <input type="text" value="NO_PARENT"/></div> <div>Fallback Enable <input type="text" value="No"/></div> <div>Fallback Cache Group <input type="text" value="EDGE-CG3"/> <input type="text" value="EDGE-CG4"/></div> <div><input type="button" value="Save"/></div>					

Step 6 Click **Save** to save the settings.

CZF File

To enable the Traffic Router to find a backup cache to use in the CZF file when the “Fallback Enable” setting of the mapped Edge cache group is set to “Yes”, you must add the longitude and latitude parameters for the backup cache groups in the CZF file.

To add the longitude and latitude settings for a cache group in the CZF file, you must add a "coordinates" section inside the cache group that contains a "longitude" and "latitude" setting, as shown in the following example:

```
{
  "coverageZones":
  {
    "Edge-West1":
    {
      "network6":
      [
        "1234:5704::/64",
        "1234:5705::/64",
        "1234:5706::/64"
      ],
      "network":
      [
        "192.168.4.0/24",
        "192.168.5.0/24",
        "192.168.6.0/24",
        "192.168.7.0/24",
        "192.168.8.0/24",
        "192.168.9.0/24"
      ],
      "coordinates":
      {
        "longitude": -118,
        "latitude": 34
      }
    },
    "Edge-West2":
    {
      "network6":
      [
        "1234:5710::/64",
        "1234:5711::/64",
        "1234:5712::/64"
      ],
      "network":
      [
        "192.168.10.0/24",
        "192.168.11.0/24",
        "192.168.12.0/24",
        "192.168.13.0/24",
        "192.168.14.0/24",
        "192.168.15.0/24"
      ],
      "coordinates":
      {
        "longitude": -122,
        "latitude": 47
      }
    },
    "Edge-West3":
    {
      "network6":
      [
        "1234:5720::/64",
        "1234:5721::/64",
        "1234:5722::/64"
      ],
      "network":
      [
        "192.168.20.0/24",
        "192.168.21.0/24",
```

```

        "192.168.22.0/24",
        "192.168.23.0/24",
        "192.168.24.0/24",
        "192.168.25.0/24"
    ],
    "coordinates":
    {
        "longitude": -117,
        "latitude": 32
    }
},
"Edge-West4":
{
    "network6":
    [
        "1234:5730::/64",
        "1234:5731::/64",
        "1234:5732::/64"
    ],
    "network":
    [
        "192.168.30.0/24",
        "192.168.31.0/24",
        "192.168.32.0/24",
        "192.168.33.0/24",
        "192.168.34.0/24",
        "192.168.35.0/24"
    ],
    "coordinates":
    {
        "longitude": -105,
        "latitude": 40
    }
}
}
}
}

```

Add a Mid Cache Group

Mid caches provide content ingest and storage functionality. When an edge cache does not contain the content requested by the client, the edge cache will proxy the request to a mid-cache server, based on the parent cache group assigned to the edge cache. If the mid-cache server does not contain the content, it is responsible for fetching the content from the Origin Server. Mid-tier caches are also organized into cache groups. Mid-tier cache groups may serve (be a parent to) multiple edge cache groups.

To add a new Mid cache group, follow these steps:



Note

For information on Origin cache groups, see [Create Origin Server Cache Groups, on page 136](#).

Procedure

- Step 1** Choose **Provisioning** > > **Edit CDN**. From the CDN Name drop-down list above the tabs, make sure the correct CDN is selected.
- Step 2** Click the **Cache Groups** tab.

Step 3 Next to the **Cache Group Name** drop-down list, click the + icon to create a new Mid cache group.

Provisioning > Edit CDN

CDN Name :

Servers Cache Groups Delivery Services Client Routing Profiles Content Invalidation

Cache Group Name

Status **success**

Short Name *

Latitude *

Longitude *

Type *

Parent Cache Group *

Secondary Parent Cache Group *

Step 4 In the New Cache Group Window that appears, enter the following information for the new group:

- **Name:** Enter a descriptive name for the group. This is the name that appears in the OMD Director drop-down menus. The name should begin with MID (in any upper case or lower case combination) so you can easily identify the type of the group.
- **Short Name:** Enter an additional descriptive name. This name is not currently used in OMD Director, but it is a required field. The Short Name can be the same as the Name.
- The **Geo Magnetic Latitude** and **Geo Magnetic Longitude:** These parameters define the geolocation of the cache group. For cache groups, the geolocation is used by the CDN system to help redirect user clients to the most optimal cache. Enter the latitude and longitude that geolocation should use for this group.

Note For more information on how the Traffic Router determines which cache group should service a request, see [Client Routing Overview, on page 155](#).

- **Type:** Choose **MID_LOC**.
- **Parent Cache Group:** If MSO is being used and specific Origin Servers should be assigned to this Mid cache group, choose the Origin Server group that contains those Origin Servers. For more information on MSO and how to configure it, see [Multi Site Origin, on page 133](#).
- **Secondary Parent Cache Group:** If MSO is being used and a different origin cache group is configured for your backup Origin Servers, choose that origin cache group as the secondary parent cache group. For more information on MSO and how to configure it, see [Multi Site Origin, on page 133](#).

Step 5 Click **Add** to add the group.

Add New Cache Group

Name *

Short Name *

Geo Magnetic Latitude *

Geo Magnetic Longitude *

Type *

Parent Cache Group *

Secondary Parent Cache Group

Step 6 When the cache group has been successfully added, the “Processing” status in the title bar will disappear and you will see a green check mark for every action in the Status column. Click **OK** to close the status window.

Note To add a server to a cache group, edit the cache server properties. For more information see [Edit an Existing Origin Server, on page 71](#).

Edit an Existing Cache Group

To edit an exiting cache group, follow these steps:

Procedure

- Step 1** Choose **Provisioning** > **Edit CDN**. From the **CDN Name** drop-down list above the tabs, make sure the correct CDN is selected.
- Step 2** Click the **Cache Group** tab.
- Step 3** From the **Cache Group Name** list, choose the group you want to edit.

Provisioning > Edit CDN

CDN Name: Delete

Servers Cache Groups Delivery Services Client Routing Profiles Content Invalidation

Cache Group Name: +

Status: success

Short Name:

Latitude:

Longitude:

Type:

Parent Cache Group:

Secondary Parent Cache Group:

Save

Step 4 If needed, change the short name, latitude, or longitude of the cache group. If you are editing an Edge or Mid cache group, you can also change the parent cache group and secondary parent cache group.

Note The Latitude and Longitude define the geolocation of the cache group. The geolocation is used by the CDN system to help redirect user clients to the most optimal cache. Enter the latitude and longitude that geolocation should use for this group.

Step 5 If you are editing an Edge cache group, you can also change the Fallback settings to configure backup edge cache groups. For more information on backup edge cache groups, see [Backup Edge Cache Groups, on page 78](#).

Note To add a server to a cache group, edit the server settings. For more information, see [Edit an Existing Origin Server, on page 71](#).

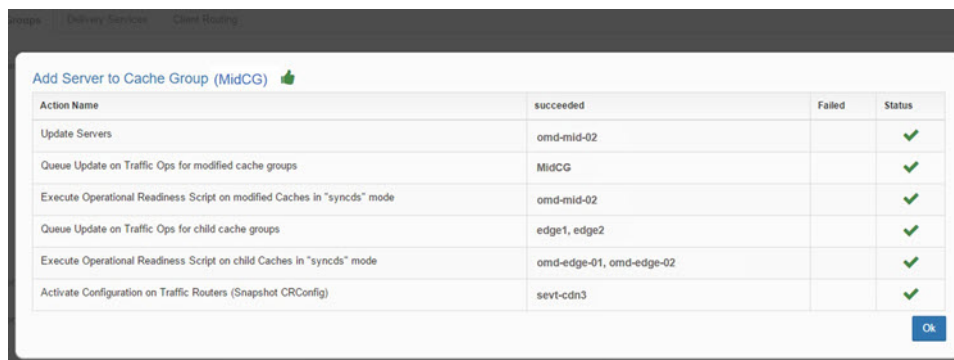
Step 6 When you have finished making your changes, click **Save**. The Processing window will appear. This window displays the progress of the cache group being updated, including the status of each action involved. For example, the following shows the process of a server being added to the cache group.

Processing ⌂
Please wait...

Add Server to Cache Group (MidCG)

Action Name	succeeded	Failed	Status
Update Servers	omd-mid-02		✓
Queue Update on Traffic Ops for modified cache groups	MidCG		✓
Execute Operational Readiness Script on modified Caches in "syncds" mode			🔄
Queue Update on Traffic Ops for child cache groups			🔄
Execute Operational Readiness Script on child Caches in "syncds" mode			🔄
Activate Configuration on Traffic Routers (Snapshot CRConfig)			🔄

Step 7 When the cache group has been successfully updated, the “Processing” status in the title bar will disappear and you will see a green check mark for every action in the Status column. The following is an example:



Step 8 Click **OK** to close the status window and return to the CDN Overview page.

Delete a Mid Cache Group



Note If you are running OMD Director in an HA configuration, you can only make configuration changes from the Primary OMD Director instance. To perform the steps in this section, make sure you are logging into the Primary OMD Director instance.

To delete a Mid cache group, follow these steps:



Note You cannot delete a Mid cache group if it is the parent of any Edge cache groups.











Procedure

Step 1 Choose **Provisioning** > > **CDN Overview**. From the CDN Name drop-down list above the tabs, make sure the correct CDN is selected.

Step 2 In the Mid Cache Groups section, click the **Delete** icon (trashcan) for the Mid cache group that you want to delete.

Note If you try to delete a Mid cache group that has servers assigned to it, you will receive an error message letting you know that the group was not deleted. You cannot delete a Mid cache group if there are any Cache servers assigned to the group.

If the word “updating” appears instead of the Delete icon (trashcan), this indicates that the Mid Cache Group has been modified and is in the process of being updated. To delete the Mid Cache Group, you will need to wait until it has finished updating.

Mid Cache Groups					
<input type="text"/>				10	
Group Name	Latitude	Longitude	Status	Action	
MID-SJC	47	-122	success	 	
MID-NTN	32.33	34.85	success	 	
MID-SG	1.28	103.85	success	 	
MID-SJ	37.23	-121.86	success	 	
MID-US-VA	36.9	-76.2	success	 	
Showing 1 to 5 of 5 entries				Previous	Next

- Step 3** In the Deleting Cache Group confirmation window that appears, click **Yes** to confirm that you want to delete the Mid cache group.
- Step 4** When the Cache Group has been successfully deleted, the “Processing” status in the title bar will disappear and you will see a green check mark for every action in the Status column. Click **OK** to close the status window.

Delete an Edge Cache Group



Note If you are running OMD Director in an HA configuration, you can only make configuration changes from the Primary OMD Director instance. To perform the steps in this section, make sure you are logging into the Primary OMD Director instance.

To delete an Edge cache group, follow these steps:

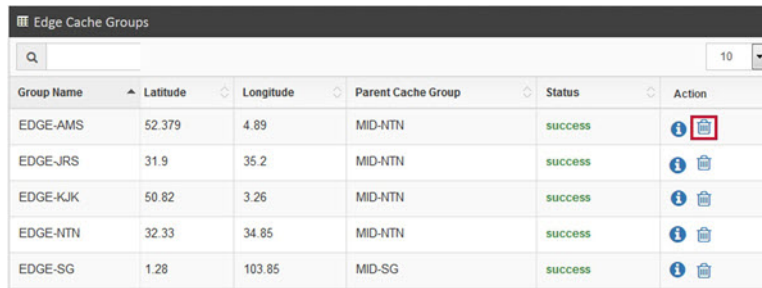
Procedure











- Step 1** Choose **Provisioning > > CDN Overview**. From the CDN Name drop-down list above the tabs, make sure the correct CDN is selected.
- Step 2** In the **Edge Cache Groups** area, click the **Delete** icon (trashcan) for the Edge cache group that you want to delete.

Note If you try to delete an Edge cache group that has servers assigned to it, you will receive an error message letting you know that the group was not deleted. You cannot delete an Edge cache group if there are any Cache servers assigned to the group.

If the word “updating” appears instead of the Delete icon (trashcan), this indicates that the Edge Cache Group has been modified and is in the process of being updated. To delete the Edge Cache Group, you will need to wait until it has finished updating.

Delete an Edge Cache Group



Group Name	Latitude	Longitude	Parent Cache Group	Status	Action
EDGE-AMS	52.379	4.89	MID-NTN	SUCCESS	 
EDGE-JRS	31.9	35.2	MID-NTN	SUCCESS	 
EDGE-KJK	50.82	3.26	MID-NTN	SUCCESS	 
EDGE-NTN	32.33	34.85	MID-NTN	SUCCESS	 
EDGE-SG	1.28	103.85	MID-SG	SUCCESS	 

Step 3 In the Deleting Cache Group confirmation window that appears, click **Yes** to confirm that you want to delete the Edge cache group.

Step 4 When the Cache Group has been successfully deleted, the “Processing” status in the title bar will disappear and you will see a green check mark for every action in the Status column. Click **OK** to close the status window.



CHAPTER 6

Manage Delivery Services

A Delivery Service is a CDN representation of an Origin Server (or a group of Origin Servers when Multi Site Origin [MSO] is configured). It defines the URL that represents the Origin Servers and which cache groups can serve content from those servers. The Delivery Service also contains the policies for serving the content to the clients, including any rewrite rules that may need to occur for the client requests or server responses.

After you have provisioned the CDN, you may need to add new or modify existing Delivery Services in the Media Streamer deployment. You may also want to configure additional features for the Delivery Services that you cannot configure using the CDN Wizard, such as URL Signing and Multi Site Origin. You can perform these tasks from the Delivery Services tab of the Edit CDN page. This chapter describes how to perform these tasks.



Note To delete a Delivery Service, you must go to Provisioning > CDN Overview.



Note If you are running OMD Director in an HA configuration, you can only make configuration changes from the Primary OMD Director instance. To perform the steps in this chapter, make sure you are logging into the Primary OMD Director instance. If you are logged into the Backup OMD Director instance, you will see “This OMD Director is currently running in backup mode” in the header. If you are logged into an OMD Director running in Detached mode because of an HA failure, you will see “This OMD Director is currently running in detached mode”.

This chapter includes the following topics:

- [Add a New DNS or HTTP Delivery Service, on page 92](#)
- [Clone a Delivery Service, on page 116](#)
- [Steering Delivery Service, on page 116](#)
- [Edit an Existing DNS or HTTP Delivery Service, on page 124](#)
- [Edit an Existing Steering Delivery Service, on page 125](#)
- [Delete a Delivery Service, on page 126](#)
- [Advanced Delivery Service Features, on page 126](#)
- [Regex Remap Settings, on page 147](#)
- [Edge Geo Blocking, on page 150](#)
- [CDN Routing Name, on page 154](#)

Add a New DNS or HTTP Delivery Service

To add a new DNS or HTTP Delivery Service, perform the following steps:

**Note**

To create a new steering Delivery Service, see [Create a Steering Delivery Service, on page 118](#).

Procedure**Step 1**

Choose **Provisioning > Edit CDN**. From the CDN Name drop-down list above the tabs, make sure the correct CDN is selected.

Step 2

Click the **Delivery Services** tab.

Step 3

Next to the **Delivery Service Name** drop-down list, click the + icon to create a Delivery Service.

Note

In addition to clicking the + icon to create a new Delivery Service, you can also click the clone icon to duplicate an existing Delivery Service. For more information see [Clone a Delivery Service, on page 116](#).

Provisioning > Edit CDN

CDN Name: [Delete](#)

Servers Cache Groups **Delivery Services** Client Routing Profiles Content Invalidation

Delivery Service Name *

Display Name *

Long Description

Use Multi Site Origin Feature ☐ Yes ☒ No

Delivery Service profile *

Routing Type * ☐ DNS ☒ HTTP ☐ STEERING ☐ CLIENT_STEERING

Content Type * ☒ Live ☐ VOD

Content Scope * ☒ National ☐ Regional

Edge cache retrieval from origin

Protocol *

Origin Server Base URL *

Customer

Content Provider

Active *

Assign Device Groups to Delivery Service

[Advanced Settings](#)

Domains for this Delivery Service:

Sub Domain * Sample Client URL

Delivery Service Server Assignments:

Assign Cache Servers	Assign Cache Servers IP	
<input type="text" value="edge-server1"/>	<input type="text" value="10.63.231.34 - Primary"/>	<input checked="" type="checkbox"/> <input type="checkbox"/>

[Cancel](#) [Save](#)

Step 4 In the **Delivery Service Name** field, enter a name for the Delivery Service.

Step 5 **Long Description:** Enter a description for the Delivery Service.

Step 6 From the **Display Name** This can be the same as the Delivery Service Name or you can enter a more descriptive display name.

Step 7 In the **Use Multi Site Origin Feature** area, click the Yes radio button to enable MSO for this Delivery Service. Otherwise leave it set to the default of No. For more information on the MSO feature, see [Multi Site Origin, on page 133](#).

Step 8 From the **Delivery Service Profile** drop-down list, choose the Delivery Service profile contains settings for advanced features of the Delivery Service, such as MSO, URL Signing, Regex Remap Parameters, and Edge Geo Blocking. If the Delivery Service you are adding will use any of these advanced features, choose the profile that contains the settings that you would like to use for these features. For more information on Delivery Service profiles, including how to create and edit them, see [Manage Profiles, on page 163](#).

Step 9 In the **Routing Type** area, choose one of the following routing types:

Note You cannot change the value of the routing type after the delivery service is created.

- **DNS:** If you choose DNS, the Traffic Router responds with an A record in response to the client request, not an HTTP 302 redirect. This A record contains a list of IP addresses for the caches in the cache group that the Traffic Router selects. Which cache the client uses from this list is based on the Proxy DNS of the client. For a DNS Delivery Service, by default the client receives a URL with **edge** prepended to the CDN domain name, for example `edge.omdcdn.example.com`. To use a custom FQDN for the delivery service, enter a different prefix to prepend in the Routing Name field of the Delivery Service.

Which IP address the Traffic Router uses to select a cache group is based on the following:

- If the Enhanced DNS Request Routing (ECS) feature is enabled in the Traffic Router profile and the ECS option is present in the DNS query, the Traffic Router will use the client subnet that is in the ECS option to select the cache group. If there are multiple ECS options in the Optional Record, the one with the longest IP prefix is used.
- If the Enhanced DNS Request Routing (ECS) feature is *not* enabled in the Traffic Router profile or if the ECS option is not present in the DNS query, the Traffic Router will use the IP address of the DNS resolver that is making the request to select the cache group.

For information on configuring the Enhanced DNS Request Routing feature, see [Configuring Enhanced DNS Request Routing, on page 271](#).

- **HTTP:** If you choose HTTP, the Traffic Router uses an HTTP 302 Redirect to control which cache the client will use. For an HTTP Delivery Service, by default the client receives a URL with `tr` prepended to the CDN domain name, for example `tr.omdcdn.example.com`. To use a custom FQDN for the delivery service, enter a different prefix to prepend in the Routing Name field of the Delivery Service.

Step 10 In the **Content Type** area, choose one of the following content types:

Note You cannot change the value of the Content Type after the Delivery Service is created.

- **Live:** Live content is not committed to disk storage. Live video segments are only buffered into the RAM disk.
- **VOD:** VOD content is cached in disk storage and the most popular objects are stored in RAM cache.
- **NO CACHE:** For a Delivery Service that has a Content Type of NO CACHE, the caches will not actually cache the content and will only be used as proxies. In this situation, the Mid tier caches are bypassed. This option is only available when HTTP is selected as the Routing Type.

Step 11 In the **Content Scope** area, choose one of the following content scopes for live content:

Note You cannot change the value of the Content Scope after the Delivery Service is created. Also, the Content Scope for VOD content is always National and you cannot change this.

- **National:** A client request will go to an Edge cache. If the content is not available on the Edge cache, the Edge cache will request the content from a Mid cache in the parent Mid cache group. If none of the Mid caches in the parent Mid cache group are available, the Edge cache requests the content from a Mid

cache in the secondary parent Mid cache group. If the Mid cache does not contain the requested content, it will request the content from the Origin Server.

If none of the Mid caches are available in either the parent or the secondary parent Mid cache group, the "Edge Cache Retrieval from Origin" setting of the Delivery Service determines whether the Edge cache can request the content directly from the Origin Server.

- **Regional:** For Delivery Services that have a Regional content scope, the Edge cache will bypass the Mid cache and request content directly from the Origin Server. This prevents Mid tier caches from having to cache regional content that it might never use.

Note Because the MSO feature uses the Mid tier caches to choose between multiple origin servers, you cannot assign a Regional content scope to a Delivery Service that has "Use Multi Site Origin Feature" set to Yes.

Step 12

The **Edge Cache Retrieval from Origin** setting only applies to Delivery Services that have a National content scope. When a Delivery Service has a National content scope, if the content the client is requesting is not available on the Edge cache and none of the Mid caches are available in either the parent or the secondary parent Mid cache group, the value of "Edge Cache Retrieval from Origin" setting determines whether the Edge cache can request the content directly from the Origin Server:

- **No:** When the "Edge Cache Retrieval from Origin" field is set to No, if none of the Mid caches are available in either the parent or the secondary parent Mid cache group, the Edge cache will *not* request content directly from the Origin Server. The request will fail with a "502 Next Hop Connection Failed".
- **Yes:** When the "Edge Cache Retrieval from Origin" field is set to Yes, if none of the Mid caches are available in either the parent or the secondary parent Mid cache group, the Edge cache will request the content directly from the Origin Server.

Note You cannot set **Edge Cache Retrieval from Origin** to Yes for a Delivery Service that has MSO enabled.

Note The **Edge Cache Retrieval from Origin** setting cannot be different for Delivery Services that use the same Origin Server.

Step 13

From the **Protocol:** drop-down list, choose from the following options:

- **HTTP:** Serves the data to the client using HTTP. Clients trying to use HTTPS will receive an SSL handshake error.
- **HTTPS:** Serves the data to the client using HTTPS. Clients trying to use HTTP will receive a 503: Service Unavailable error.
- **HTTP and HTTPS:** Data will be served to the client using either HTTP or HTTPS, depending on what the client requests.
- **HTTP to HTTPS:** Non-secure (HTTP) clients are redirected with a 302 redirect message to use a secure (HTTPS) connection. Secure (HTTPS) clients will continue with their HTTPS connection.

If you choose any of the options that contain HTTPS, you must also provide the SSL details. Follow these steps to enter the SSL details:

- a) Click the shield icon.

Provisioning > Edit CDN

CDN Name: [Delete](#)

Servers Cache Groups **Delivery Services** Client Routing Profiles Content Invalidation

Delivery Service Name *

Display Name *

Long Description

Use Multi Site Origin Feature * ☐ Yes ☒ No

Delivery Service profile *

Routing Type * ☐ DNS ☒ HTTP ☐ STEERING ☐ CLIENT_STEERING

Content Type * ☒ Live ☐ VOD

Content Scope * ☒ National ☐ Regional

Edge cache retrieval from origin

Protocol * [Click to enter SSL details](#)

Origin Server Base URL *

- b) The SSL Certificate Details window appears. For SSL Mode you can generate an SSL certificate or you can add an existing certificate for the Delivery Service.

- **Generate:** If you choose this option, enter all the required values, indicated by an asterisk (*), in the SSL Certificates Detail window and then click **OK**:

SSL Certificate details

SSL Mode * ☒ Generate ☐ Add

Country Code *

State Name *

City Name *

Organization Name *

Organizational Unit Name *

[Ok](#) [Cancel](#)

- **Add:** If you choose this option, enter all the required values, indicated by an asterisk (*), in the SSL Certificates Detail window and then click **OK**:



The image shows a dialog box titled "SSL Certificate details" with a close button (X) in the top right corner. Inside the dialog, there are three sections, each with a label and a text input field:

- SSL Mode ***: Two radio buttons are present. The first is labeled "Generate" and is unselected. The second is labeled "Add" and is selected (indicated by a blue dot).
- Private Key ***: A text input field with the placeholder text "Enter Private Key".
- Certificate Signing Request (CSR)**: A text input field with the placeholder text "Certificate Signing Request".
- Certificate (CRT)***: A text input field with the placeholder text "Enter Certificate".

At the bottom right of the dialog, there are two buttons: "Ok" (in blue) and "Cancel" (in grey).

Step 14

In the **Origin Server Base URL** field, enter the URL used to reach the Origin Servers that host the content for this Delivery Service. The value in this field based on the following criteria:

- If this Delivery Service is not using MSO and the Origin Server for this Delivery Service is not used by any other Delivery Services, enter the URL of the Origin Server.
- If this Delivery Service is not using MSO but the Origin Server for this Delivery Service will be used by another Delivery Service, enter a URL where the FQDN portion of the URL is different from the actual FQDN that is used to reach the Origin Server. You will also need to add the Origin Server to the

CDN and assign the Origin Server to this Delivery Service from either the Assign Cache Servers list or by choosing a device group that contains the Origin Server from the Assign Device Groups to Delivery Service list. For information on how to add the Origin Server to the CDN, see [Add an Origin Server to the CDN, on page 63](#).

- If this Delivery Service is using MSO, enter a URL where the FQDN portion of the URL is different from the FQDN of *any* of the Origin Servers that this Delivery Service will use.

Note To support MSO and device groups, every Delivery Service within a CDN must have a unique Origin Server Base URL (also referred to as the OFQDN), with the following exception: If two Delivery Services use the exact same MSO settings and use the exact same Edge caches, Mid caches, Origin Servers, and Device Groups, then two Delivery Services can use the same OFQDN.

Step 15 The **Customer** field is for backwards compatibility only. For OMD Director 3.11 and later, use the **Content Provider** field instead.

Step 16 From the **Content Provider** drop-down list, choose the content provider organization to associate with this delivery service. This is an optional field, however, the content provider field plays an important role in controlling who can view the OMD Insights analytics information for this Delivery Service. Choose one of the following:

- **Content Provider Viewers and Content Provider Admins:** OMD Director users that have the role of Content Provider Viewer or Content Provider Admin can only view OMD Insights information for Delivery Services that are assigned a Content Provider that matches the Content Provider organization they are assigned.
- **Reseller Viewers and Reseller Admins:** OMD Director users that have the role of Reseller Viewer or Reseller Admin can only view OMD Insights information for Delivery Services that are assigned a Content Provider that is assigned to the Reseller organization they are assigned.

For example, if user jsmith has the role of Content Provider Admin and has been assigned the Organization of CompanyX, and a Delivery Service named VOD-SJ has been assigned a Content Provider of CompanyY, jsmith *cannot* view any of the OMD Insights information for the VOD-SJ Delivery Service.

Note The Content Provider field only affects users that have one of the Content Provider or Reseller roles. It does not limit or restrict users that have been assigned any other role.

Step 17 From the **Active** drop-down list, choose **Yes** to enable the Delivery Service.

Step 18 In the **Assign Device Groups to Delivery Service** field, choose the device groups that contain the cache servers and Origin Servers that you would like to assign to the Delivery Service. For more information on device groups, see [Device Groups Overview, on page 181](#).

Note If you assign servers (Edge cache, Mid cache, or Origin) from both the Assign Cache Server list and the Assign Device Groups to the Delivery Service list, the servers that serve the Delivery Service will be a union of the two. When the MSO algorithm assigned to the Delivery Service is False, all Mid cache servers that are in this union that are not assigned a Parent Cache Group will use the rank of any Origin Servers that are in this union to determine the primary and backup Origin Servers. The Origin Server with the lowest rank is considered the primary. Also, all of the Mid caches that are in the parent cache groups of any specifically assigned Edge cache (from the “Assign Cache Servers” field) are automatically assigned to the Delivery Service.

- **Routing Name:** By default, a client receives a URL with edge prepended to the CDN domain name for Delivery Services that have a Routing Type of DNS and "tr" prepended to the CDN domain name for Delivery Services that have a Routing Type of HTTP. To use a prefix other than the default, enter a different prefix in the Routing Name field.

- **Sub Domain:** You can create the sub domain for a Delivery Service in two different ways:
 - **Enter only the sub domain:** Enter only the sub domain portion of the FQDN. This subdomain is prepended to the CDN domain and either "tr" (for an HTTP Delivery Service) or "edge" (for DNS Delivery Service) is the hostname used to create the URL for client requests to this Delivery Service. For example, if the CDN domain name is spcdn.company.example and you enter a sub domain of on-site, the client will use a URL for this sub domain of tr.on-site.spcdn.company.example for an HTTP Delivery Service and a URL of edge.on-site.spcdn.company.example for a DNS Delivery Service. The following is an example:

Domains for this Delivery Service:

Sub Domain * Sample Client URL

on-site http://tr.on-site.spcdn.company.example

- **Enter a full FQDN:** A custom FQDN enables you to use a different hostname than the default hostname of "tr" for an HTTP Delivery Service or "edge" for the DNS Delivery Service. For this option you enter the full FQDN: hostname, sub domain you want to create, and domain of the CDN. for example srv1.on-site.spcdn.company.example. The following is an example:

Note You can only add one custom FQDN to a Delivery Service.

Domains for this Delivery Service:

Sub Domain * Sample Client URL

srv1.on-site.spcdn.company.example http://srv1.on-site.spcdn.company.example

Step 19 In the **Assign Cache Servers** area, choose one of the following:

- Choose an Edge cache server that should provide cache and streaming services for this Delivery Service. You do not need to configure any Mid cache servers because they are inherently assigned through the Edge cache parent relationship.

To add additional Edge cache servers, click the + icon.

- If the Delivery Service is using MSO, or if you have an Origin Server that is shared by more than one Delivery Service, choose the Origin Servers that will host content for this Delivery Service. Optionally, you can choose a device group from the Assign Device Groups to Delivery Service list that contains the Origin Servers that will host content for this Delivery Service.

Step 20 In the **Assign Cache Servers IP** area, for each Edge cache server you assign to the Delivery Service, choose which IP address it should use to stream content for this Delivery Service.

Note By default, only the primary IP address that is assigned to the server in Linux will be available. To use a secondary IP address that is assigned to the server in Linux, see [Manage Secondary Streaming IPs](#), on page 68.

Step 21 To configure the advanced settings for the Delivery Service, click the **Advanced Settings** link below the Assign Device Groups to Delivery Service field. The Additional Info window appears. See the following sections for information on the fields in the Additional Info window.

Step 22 When you are finished entering the information for the Delivery Service, click **Save**.

The Processing window appears. This window displays the progress of the Delivery Service being created, including the status of each action involved.

Delivery Service create (http-mos1) Show/Hide Config

Action Name	succeeded	Failed and Reasons	Status
Create delivery_service http-mos1 in Traffic Ops	http-mos1		✓
Update cache servers ATS version into profile for delivery_service http-mos1	mid-cache2, edge-cache5		✓
Queue Update on Traffic Ops for mid cache groups	MID		✓
Execute Operational Readiness Script on mid Caches in "syncds" mode	mid-cache2		✓
Queue Update on Traffic Ops for edge cache groups	EDGE		✓
Execute Operational Readiness Script on edge Caches in "syncds" mode			↻
Activate Configuration on Traffic Routers (Snapshot CRConfig)			↻

When the system has successfully created the Delivery Service, the "Processing" status in the title bar will disappear and you will see a green check mark for every action in the Status column.

Delivery Service create (http-mos1) ✓ Show/Hide Config

Action Name	succeeded	Failed and Reasons	Status
Create delivery_service http-mos1 in Traffic Ops	http-mos1		✓
Update cache servers ATS version into profile for delivery_service http-mos1	mid-cache2, edge-cache5		✓
Queue Update on Traffic Ops for mid cache groups	MID		✓
Execute Operational Readiness Script on mid Caches in "syncds" mode	mid-cache2		✓
Queue Update on Traffic Ops for edge cache groups	EDGE		✓
Execute Operational Readiness Script on edge Caches in "syncds" mode	edge-cache5		✓
Activate Configuration on Traffic Routers (Snapshot CRConfig)	CDN1		✓

OK

Step 23

Click **OK** to close the status window and return to the Edit CDN page.

Step 24

To confirm that the Delivery Service is saved, ensure that the new Delivery Service appears in the **Delivery Service Name** drop-down list.

Delivery Service Advanced Settings

To configure the advanced settings for the Delivery Service, click the **Advanced Settings** link below the Assign Device Groups to Delivery Service field. The Additional Info window appears. This window is divided into four sections: General, Content Preposition, Header Rewrite Configurations, and URL Signing. Use the following information to configure the advanced settings.

Advanced Settings: General

Additional Info

General:

Query String Handling	0 - use qstring in cach	Initial Dispersion	1
Regex Remap Expression		IPv6 Routing Enabled?	No
Bypass FQDN		Range Request Handling	0 - Don't cache
Maximum Bits per Second allowed globally (4T or 500M are valid entries)		Delivery Service DNS TTL	3600
Maximum Transaction allowed globally		Geo Miss Default Latitude	41.881944
Traffic Router Log Request Headers		Geo Miss Default Longitude	-87.627778
Pacing Rate	0 bps		

- Query String Handling:** Query strings are the set of key/value pairs that occur after the ? in a URL. The Query String Handling setting enables you to control how requests that contain query strings are handled by the CDN. The options for this setting are:
 - 0 - use qstring in cache key, and pass up:** With this setting, the query string is saved as part of the URL cache key on the cache and will be used to identify content. Therefore, two URLs that are the same except for the query string will be saved as separate URL cache keys. In addition, the query string is passed up to either the Mid cache or Origin Server. This option is good to use when the content that is returned by the Origin Server depends on the query string. This is the default setting.
 - 1 - ignore in cache key, and pass up:** With this setting, the query string is not saved as part of the URL cache key. Therefore, two URLs that are the same except for the query string will match the same URL cache key. In addition, the query string is passed up to either the Mid cache or Origin Server. This option is good to use when the content that is returned does not depend on the query string, but the Origin Server still requires the query string for other purposes, such as authentication or logging.
 - 2 - drop at edge:** With this setting, the query string is not saved as part of the URL cache key. Therefore, two URLs that are the same except for the query string will match the same URL cache key. With this setting, the query string is not passed up to either the Mid cache or Origin Server.
- Regex Remap Expression:** This setting enables you to configure a regular expression remap rule that modifies the URL that the client is requesting. The base remap rule modifies the original client request for retrieving content from the origin listed in the Base Origin Server URL field. Placing a rule in this field can override and replace the request URL. If the regular expression does not match the request URL, the Base Origin Server URL is used.



Note This feature cannot be used when the Query String Handling setting is set to “2 - drop at edge”.

The following is the syntax for this field:

regex_expression_fromURLpath remap_toURL options

- *regex_expression_fromURLpath*: A regular expression that is evaluated to determine which client requests to remap. Any client requests that match this regular expression are remapped. By default, the regular expression only matches against the URL path and query string. The path will always start with a “/”.
- *remap_toURL*: The URL to which you want to remap the matching client requests. Various substitution strings are allowed in the remap_toURL parameter during evaluation. For a list of substitution strings that can be used in this parameter, please refer to the open source documentation at https://docs.trafficserver.apache.org/en/5.3.x/reference/plugins/regex_remap.en.html.
- *options*: Additional options that you want to apply to the remap rule. For a list of available options, please refer to the open source documentation at https://docs.trafficserver.apache.org/en/5.3.x/reference/plugins/regex_remap.en.html.

For example:

```
^(vod.*)/more http://www.vod.example/$h/$0/$1
```



Note By default, only the path and query string of the URL are provided for the regular expressions to match. The Regex Remap Settings tab of the Delivery Service profile enables you to specify parameters and settings that you can use to modify the behavior of what regular expressions will match. For more information on how to configure the Regex Remap parameters, see [Regex Remap Settings, on page 147](#).

- **Bypass FQDN**: The Bypass FQDN is used in the following situations:
 - **Thresholds exceeded**: If either the “Maximum Bits per Second Allowed Globally” or the “Maximum Transaction Allowed Globally” limits are reached, the Traffic Router will send any overflow traffic to the FQDN listed in the Bypass FQDN field. To redirect the requests to the Origin Server, you would enter the FQDN of the Origin Server. By configuring Bypass FQDN, the Maximum Bits per Second Allowed Globally, and the Maximum Transaction Allowed Globally settings, you can essentially place a quota on how much traffic and how many computing resources the Delivery Service can use.



Note You do not have to configure both the Maximum Bits per Second Allowed Globally and the Maximum Transaction Allowed Globally settings. However, you need to configure at least one of them for the Bypass FQDN setting to have any affect.

- **No caches available:** If there are no caches available when the Traffic Router tries to route the client request, the Traffic Router will redirect the client request to the FQDN listed in the Bypass FQDN field. To redirect the requests to the Origin Server, you would enter the FQDN of the Origin Server. This essentially works like a “Last Resort Routing”.
- **Maximum Bits per Second Allowed Globally:** This setting is used in conjunction with the Bypass FQDN setting. The value that you enter in this field determines the maximum bits per second (bps) this Delivery Service can serve across all Edge caches. When this limit is reached, the Traffic Router diverts the overflow traffic to the Bypass FQDN destination.



Note For this setting to have any effect, you must also configure the Bypass FQDN setting.

- **Maximum Transaction Allowed Globally:** This setting is used in conjunction with the Bypass FQDN setting. The value that you enter in this field determines the maximum transactions per second that this Delivery Service can serve across all Edge caches. When this limit is reached, the Traffic Router diverts the overflow traffic to the Bypass FQDN destination.



Note For this setting to have any effect, you must also configure the Bypass FQDN setting.

- **Traffic Router Log Request Headers:** This field enables you to specify headers that you would like to include in the Traffic Router access log entries. To specify which headers to include, enter a list of header keys separated by `__RETURN__`. For example, "Header1-Name `__RETURN__` Header2-Name". You can add additional headers by separating them with "`__RETURN__`". These headers will appear after the "rh=" token. If the listed header name is not present, it will not be logged.



Note You must have completed the upgrade steps in the *Cisco Media Streamer and Cisco Media Broadcaster Installation and Upgrade Guide* for the Cache servers before you will be able to set this value.

- **Pacing Rate:** Enter a maximum bitrate per session (TCP connection) to use for delivering content.
- **Initial Dispersion:** The Initial Dispersion setting determines how many copies of requested content are stored in the cache group. By default, the Initial Dispersion setting is 1, which means that only a single copy of the requested content is cached in the cache group. Therefore, by default only one cache server will have a copy of the content. If the cache server that contains the content goes offline or is overloaded, the content is moved to another cache server in the cache group. This optimizes the amount of storage that is available for caching content.

However, if you have a piece of content that you know is going to be very popular, you may want OMD to cache more than one copy of the content. These copies would be created at the time of the initial request and the system would not have to wait for load or cache health settings to be reached before additional copies of the content are created. This would result in faster performance for the clients and better control of the bandwidth utilization. However, keeping multiple copies of the same content will reduce the amount of storage available for caching unique content.

To configure OMD to cache more than one copy of the content that is served by the Delivery Service, from the Initial Dispersion drop-down list choose the number of copies of the content that you want to cache. These copies would be stored on different cache servers in the cache group. For example, if you choose 3, a copy of the content is stored on 3 different cache servers.

- **IPv6 Routing Enabled?:** This setting determines whether the Traffic Router will respond to IPv6 DNS (AAAA record) requests for the hostnames of the Delivery Service. The default setting is No, which means the Traffic Router will only respond to IPv4 DNS (A record) requests for the hostnames of the Delivery Service. If you change this setting to Yes, the Traffic Router will respond to both IPv6 DNS (AAAA record) requests and IPv4 DNS (A record) requests for the hostnames of the Delivery Service.
- **Range Request Handling:** Some of the content that the CDN may deliver can be very large, however the client may only need a portion of that file because they are only watching a portion of the content.

To enable a client to request only a portion of a file, browsers and video players support a Range Request header, which denotes what part (bytes) of the content the client is requesting. For example, a client could make a request that they only want bytes 0 through 199 of the content (the first 200 bytes) or they may request that they want bytes from a range later in the file such as bytes 10,000 and 11,000. The Range Request Handling setting determines how the cache server handles the content from these requests. The options for this setting are:

- **0- Don't Cache:** When you choose this option, the cache servers will not cache the content for any requests that contain a range request (essentially any partial object request). The cache servers pass through the request in order to delivery the content to the client, but they do not cache the content that they receive from the Mid cache group or Origin Servers. This is the default setting.
- **1 - Use background_fetch plugin:** When you choose this option, if a client makes a range request, for example they request only bytes 200 - 300 of an object, the cache server will cache the entire object but will return to the client only the requested range. The cache server delivers the range to the client as soon as the range is downloaded; it does not have to wait for the entire object to be cached before it can deliver the requested range to the client. However, the cache server will save the entire object in cache, in the background, in case another client wants either a portion of this object or the whole object.
- **2 - Use cache_range_requests plugin:** When you choose this option, the cache server will store requests for the same object but different byte ranges as separate unique objects in cache. For example, if a client asks for a byte range of 0 - 10 of a certain object, the cache server will cache those bytes as one object. If another client asks for a different byte range for the same object, for example bytes 5 - 10, the cache server will store this request as a completely separate unique object in cache, even though the byte range overlaps with a previous request.



Note

If you have clients that make range requests and you can control these requests so that they are aligned (the clients are requesting the same ranges for the same object), choosing option “2 - Use cache_range_requests plugin” would be beneficial. However, if you have many clients that make range requests but those range requests are not aligned, option 2 is not efficient because having a large number of range requests that are not aligned can result in a large number of objects with overlapping content stored in cache

- **Delivery Service DNS TTL:** This setting determines the TTL, in seconds, that the Traffic Router will set on the A record responses to the clients for the hostnames of this Delivery Service. When the IPv6

Routing Enabled setting is set to Yes, this TTL also applies to the AAAA record responses. The default is 3600 seconds.

- **Geo Miss Default Latitude and Geo Miss Default Longitude:** What these settings determine, depends on the value of the Geo Limit setting. For more information on the Geo Limit setting, see [Advanced Settings: Geo Limit, on page 113](#).

- **If the Geo Limit field is set to None:** If the Geo Limit is set to None, the Traffic Router will allow all client requests through and will use the CZF file to determine which cache group to direct the client to. If the check of the CZF file does not return a Cache to use for the request, Proximity routing is checked next, if it is enabled in the profile of the Traffic Router. (See [Proximity Routing, on page 159](#) for more details.)

If Proximity routing is not enabled, or if Proximity routing does not return a Cache to use for the request, the Traffic Router will use geolocation based routing to determine a cache group the client request should use and then chooses a cache from that group. If no geolocation is available for the requesting address, the Geo Miss Default Latitude and Geo Miss Default Longitude settings of the Delivery Service are used to determine the closest cache group to use for the client request and then chooses a cache from that cache group. (See [Geolocation Based Routing, on page 161](#) for more details.)

- **If the Geo Limit field is set to CZF Only:** If the Geo Limit is set to CZF file Only, the Geo Miss Default Latitude and Geo Miss Default Longitude settings are not used. This is because with this setting, only clients whose IP addresses match an entry in the CZF file are allowed through. If the client matches an entry in the CZF file, that entry will determine which cache group the client is directed to.
- **If the Geo Limit field is set to CZF + CountryCode(s):** If Geo Limit is set to CZF file + CountryCode(s), if the client IP address does not match an entry in the CZF file, then the Traffic Router tries to determine the location of the client IP address using a geolocation database. If the Traffic Router cannot find the location of the client using the geolocation database, it will then use the Geo Miss Default Latitude and Geo Miss Default Longitude fields to determine what country the request is from.

If the country location of the client matches a country listed in the Geo Limit Country Codes field, the Traffic Router will then check the client request against the Anonymous Blocking configuration to determine whether the request is allowed. If the client request is allowed, the Traffic Router uses Proximity routing, if it is enabled, to determine which cache group to use. (See [Proximity Routing, on page 159](#) for more details.)

If Proximity routing is not enabled, or if Proximity routing does not return a Cache to use for the request, the Traffic Router will use the Geo Miss Default Latitude and Geo Miss Default Longitude settings of the Delivery Service to determine the closest cache group to use for the client request. (See [Geolocation Based Routing, on page 161](#) for more details.)

Advanced Settings: Content Preposition

Content Preposition: ▲

DSCP Edge Tag	<input type="text" value="0 - Best Effort"/>	Session Tracking	<input type="text" value="Yes"/>
DSCP Mid Tag	<input type="text" value="--"/>	Session Tracking Query Key List	<input type="text" value="sid, uid"/>
Use Content Prepositioning	<input type="text" value="Yes"/>		
Ingest Manifest File URL *	<input type="text" value="http://store.cdn.company"/>		
HTTP Proxy for Ingest Manifest File URL	<input type="text" value="http://proxy.company.exz"/>		
Preload on	<input type="text" value="edge only"/>		
Use Dedicated Volume	<input type="text" value="hostname=s1.company.1"/>		
Cache Configuration	<input type="text" value="dest_host=os1.company"/>		

- **DSCP Edge Tag:** All of the packets that are sent from the Edge cache to the client can be marked with a DSCP value. The DSCP Edge Tag setting enables you to configure what that DSCP value will be. The DSCP Edge Tag can be used by routers and other network devices to give the traffic from this Delivery Service different priority.
- **DSCP Mid Tag:** All of the packets that are sent from the Mid cache to the Edge cache can be marked with a DSCP value. The DSCP Mid Tag setting enables you to configure what that DSCP value will be. This tag can be used by routers and other network devices to give the traffic from this Delivery Service different priority.
- **Use Content Prepositioning:** Typically content is not acquired and stored on the cache until a client requests the content. However, in some situations, for example if you have content that you know is going to be very popular, you may want the cache to acquire and store the content before it is requested. This is what the Content Prepositioning feature enables you to do.

To configure the Delivery Service to support content prepositioning, choose Yes for the Use Content Prepositioning setting. In addition, you must also configure the Ingest Manifest File URL field and the Preload On field for content prepositioning to occur. These fields are discussed below. The default for this setting is No.



Note

The following fields are only available to edit if Use Content Prepositioning is set to **Yes**.

- **Ingest Manifest File URL:** When you configure prepositioning, you use an Ingest Manifest file to define what content to put on the selected caches ahead of time. In the Ingest Manifest File URL field, enter the web address (URL) for the location of the Ingest Manifest file. For example, `http://store.company.example/config/cdnmanifest.json`.



Note For an example of an Ingest Manifest file and for a description of its syntax, please refer to [Example Ingest Manifest File, on page 261](#).

- **HTTP Proxy for Ingest Manifest File URL:** If the cache servers need to go through a proxy server to retrieve the Ingest Manifest file configured in the Ingest Manifest File URL field, enter the URL for the proxy server in this field.
- **Preload on:** This field enables you to choose the type of cache servers that you would like preposition the content on. Your options are Edge Only, Mid Only, or Edge + Mid.



Note The following two fields are optional fields that you can configure for Content Prepositioning. These features require additional configuration outside of OMD Director in Traffic Ops.

- **Use Dedicated Volume:** By default cache servers have two volumes:
 - One volume for live content for Live Delivery Services, which is stored in RAM
 - One volume that is typically all of the disks in the cache server, which is used for VOD Delivery Services

Additionally, a third volume can be created that is dedicated to storing prepositioned content for a Delivery Service.

If a third volume has been created on the cache server, the Use Dedicated Volume field points to that volume, which tells the cache server to use that volume for prepositioned content. With this field you also configure which Origin Servers in the Delivery Service you want to preposition content from. The following is the syntax for this field:

hostname=server.company.example volume=X

where *server.company.example* is the URL of the Origin Server whose content you want to preposition on this volume and *X* is the volume number that was created for the dedicated volume. This number must match the Preposition Volume number that was created in Traffic Ops. For example **hostname=s1.company.example volume=3**.

- **Cache Configuration:** By default when content is prepositioned on the cache server, the content stays on the cache until the cache server needs space to hold new content. When the cache fills up and clients make new requests to the cache, the cache automatically evicts the oldest content, which is not always the behavior that you want. When you preposition content in the cache, you might want to ensure that the prepositioned content is always available and that it is never evicted from the cache. The pinning feature enables you to accomplish this.

If the cache server has been configured to support pinning in Traffic Ops, you use the Cache Configuration field to configure which Origin Servers you will pin content for and for how long the content will be pinned. The following is the syntax for this field:

dest_host =server.company.example pin-in-cache=time

where *server.company.example* is the URL of the Origin Server whose content you want to pin and *time* is the length of time to pin this content. For example **dest_host=os1.compay.example**

pin-in-cache=100d. In this example 100d means to pin the content for 100 days. For pin-in-cache you can reference “d” for day, “h” for hours, and “m” for minute. If the hostname in the Cache Configuration field matches the hostname in the Use Dedicated Volume field, content for this Origin Server will be pinned to the volume configured in the Use Dedicated Volume field.

- **Session Tracking:** Session tracking assigns all CDN viewers tracking values within an HTTP cookie. Those tracking values are written into the OMD Transaction Logs for further analysis.

To enable session tracking, from the Session Tracking drop-down list choose **Yes**.



Note Session tracking is currently only supported on Edge caches and Mid caches. Traffic Router does not support session tracking.

Session tracking occurs before header rewrite. When any header rewrite rules are defined to rewrite the “Cookie”/“Set-Cookie” field, carefully define the rule so that it does not interfere with session tracking or the Session Tracking Security Enablement feature.

- **Session Tracking Query Key List:** Session tracking can use either a randomly generated ID or can accept customer parameters from a URL query string. To use customer parameters, in the Session Tracking Query Key List field, enter a comma separated list of URL Query Parameters to capture in the cookie.
- **Session Tracking Security Enablement:** To enable the Session Tracking Security Enablement feature, in the **Session Tracking Query Key List** field enter **key=<keyvalue>**, where <keyvalue> is the secret key to use to create the cookie signature. If URL Query Parameters also need to be listed in the Session Tracking Query Key List field for session tracking, add the **key=<keyvalue>** at the end separated by a comma. For example: **sessionId, key=mykey**.



Note To use the Session Tracking Security Enablement feature, Session Tracking must be set to Yes and you must have the URL Signing feature configured. For more information on configuring the URL Signing feature, see [URL Signing, on page 127](#).

For more information on Session Tracking Security Enablement, see [Session Tracking Security Enablement, on page 132](#).

Advanced Settings: Header Rewrite Configurations

In addition to remapping a client request to an origin server, there may be a need to modify information in either the client request or the response header. The Header Rewrite feature enables you to define rules to modify the headers for both requests and responses. The rules that you define will determine which requests or responses will be modified, based on matching conditions, and how the header will be modified, based on the action (operator) you define.



Note These header rewrite rules are configured on the Delivery Service and will apply to all of the Edge caches, Mid caches, or Traffic Routers that service that Delivery Service depending on the type of header rewrite rules you define.

Header Rewrite Configurations :

Edge Header Rewrite Rules

```
cond %{READ_RESPONSE_HDR_HOOK}
rm-header Set-Cookie
```

Traffic Router Additional Response Headers

Cache-C

no-cache

Mid Header Rewrite Rules

Cache URL expression

- **Edge Header Rewrite Rules:** The Edge cache will modify either the request or response header, based on the rule you define. Enter the rewrite rules in the text box. To enter more than one rewrite rule, separate them by commas. See [Header Rewrite Rules Syntax, on page 267](#) for information on the syntax for this rule.
- **Mid Header Rewrite Rules:** The Mid cache will modify either the request or response header, based on the rule you define. To enter more than one rewrite rule, separate them by commas. See [Header Rewrite Rules Syntax, on page 267](#) for information on the syntax for this rule.



Note Mid Header Rewrites rules are applied based on the Origin FQDN. If multiple Delivery Services share the same origin, mid header rewrites may conflict with each other depending on the contents.

- **Cache URL Expression:** The Cache URL Expression field enables you to change the cache key that is used for caching a request. This enables different URLs that provide the same content to use the same cached object. The format for the entries in this field is:

<pattern> <replacement>

- *<pattern>* is a regular expression that will be applied against the incoming request URL to determine which URLs to match.

- *<replacement>* is the cache key to use for the incoming request URLs that match *<pattern>*.

Example:

```
http://s[123].example.com/(.*) http://s.example.com.TSINTERNAL/$1
```

With this example, the s1.example.com, s2.example.com, and s3.example.com domains will effectively share the same cache objects.



Note To enter more than one URL cache expression, enter each expression on a separate line.

- **Traffic Router Additional Response Headers:** The Traffic Router will add additional HTTP headers in the response to the client, based on the information you enter. In the Name field, enter the name of an additional header that you want the Traffic Routers to add and in the Value field, enter the value to assign to the header. You can enter multiple headers by clicking the + button to add a new Name/Value pair.



Note This settings is only supported on HTTP-based Delivery Services. You will not see it if you are configuring the Advanced Settings of a DNS-based Delivery Service.



Note Its preferred to configure any CORS related headers using the **CORS** section, instead of placing them here.

Advanced Settings: URL Signing

The URL signing feature of Media Streamer provides the infrastructure to validate content URLs to prevent unauthorized access.

URL Signing :

Signing Algorithm	url_validator
Signing Keys From	upload key
URL Signing Keys	<pre>"SYM-KO1-KN2": "FMwlx_hFBe_kZRVGuFj8Q1VGfAuxnyuY", "AES-KO1-KN2": "cPmyPmnwDXTMuTy7", "ASYM-KO1-KN2": "CoUKbYsyG9A3jZbjP30IEVbt)EojOq", "ASYM-KO10-KN11": "Mkz3hgV10J89LGUxiMulhJFIq73Y0Dz"</pre>

- **Signing Algorithm:** To use the URL signing feature, from the Signing Algorithm drop-down list, choose **url_validator**. After you choose url_validator, configure the following fields:

- **Signing Keys From:** The first time you configure the URL signing feature for the Delivery Service, you must identify keys for the Delivery Service to use. Any Edge cache servers that serve content for this Delivery Service will use these keys to validate the content requests. To identify the keys to use, do one of the following:
 - **upload key:** To upload specific keys for this Delivery Service to use, choose **upload key** from the "Signing Keys From" drop-down list. You must then enter the keys in the "URL Signing Keys" field.



Note If this is the first time you are configuring URL signing for the Delivery Service, do not choose static key. The static key option is only for when you are editing a Delivery Service.

- **URL Signing Keys:** If you chose upload key from the "Signing Keys From" drop-down list, enter the keys to upload for this Delivery Service in this field. The following is the format to enter the keys. Make sure to separate the key entries with a comma:

`"KEY_Type-KO#-KN#": "key"`

where:

- **KEY_Type:** Identifies the type of key: SYM, AES, or ASYM.
- **KO#:** Identifies the key owner.
- **KN#:** Identifies the key number.
- **key:** This is the actual key.

Optionally, you can include an `error_url` parameter that determines what HTTP response the Edge cache returns to the client if the URL validation fails. If you do enter an `error_url` parameter, the default response the client receives is a “403: Forbidden” response. The following is the format for the error URL parameter:

```
"error_url": "http_response_code"
```

If the `http_response_code` is a Redirection (3XX) code, you must also specify the URL to which the client should be redirected. For example, `"error_url": "302 http://abc.com/url_signing_fail/"`

The following is an example of an entry for the **URL Signing Keys** field:

```
"SYM-KO1-KN2": "FMwIx_hFBe_kZRVGuFj8Q1VGfAuxnyuY",
"AES-KO1-KN2": "cPmyPmnwDXTMuTy7",
"ASYM-KO1-KN2": "CoUKbYsyeaG9A3JZbjP3oIEVbtOEcjOq",
"ASYM-KO10-KN11": "Mcz3hqV10J89LGUxiMulhJlFlg72Y0dz"
"error_url": "302 http://abc.com/url_signing_fail/"
```

Advanced Settings: Traffic Router CORS Settings for HTTP(s) Content Routing

CORS enables browsers to use a predefined set of headers and methods to describe the set of origins that are permitted to access resources using a web browser. “Origin” as it pertains to CORS and the CORS configuration refers to the scope of authority or privilege used by user agents (for example, web browser) as specified in an “Origin:” header in the requests. It does not refer to the Origin Server of the CDN.

Configure the following settings to enable and configure the CORS policies that need to be enforced in your CDN. (Media Streamer supports both simple and non-simple (preflight) CORS requests.)



Note

The CORS settings are only supported on HTTP-based Delivery Services. You will not see them if you are configuring the Advanced Settings of a DNS-based Delivery Service.

Traffic Router CORS Settings for HTTP(s) Content Routing :

<p>Allowed Origins</p> <div style="border: 1px solid #ccc; padding: 2px; margin-bottom: 5px;">https://example1.company</div> <div style="text-align: center;"> - + </div>	<p>Preflight Allowed Methods</p> <div style="display: flex; gap: 10px;"> <div><input checked="" type="checkbox"/> OPTIONS</div> <div><input checked="" type="checkbox"/> PUT</div> </div> <div style="display: flex; gap: 10px;"> <div><input type="checkbox"/> DELETE</div> <div><input type="checkbox"/> TRACE</div> </div> <div><input type="checkbox"/> CONNECT</div>	<p>Allow Credentials</p> <div style="border: 1px solid #ccc; padding: 2px; margin-bottom: 5px;">Yes</div>	<p>Preflight Allowed Headers</p> <div style="border: 1px solid #ccc; padding: 2px; margin-bottom: 5px;">X-Requested-With</div> <div style="border: 1px solid #ccc; padding: 2px; margin-bottom: 5px;">Client-Security-Token</div> <div style="text-align: center;"> - + </div>
<p>Exposed Headers</p> <div style="border: 1px solid #ccc; padding: 2px; margin-bottom: 5px;">Content-Length</div> <div style="border: 1px solid #ccc; padding: 2px; margin-bottom: 5px;">Accept-Encoding</div> <div style="text-align: center;"> - + </div>	<p>Preflight Max Age</p> <div style="border: 1px solid #ccc; padding: 2px; margin-bottom: 5px;">300</div> <div style="background-color: #d9ead3; padding: 2px; margin-bottom: 5px;">Sec</div>		

Save

- **Allowed Origins:** Enter a URI that may access the resources of the Delivery Service. For example, `cdn.example1.com` or `https://example2.com:8080`. The Traffic Router compares the value in the Origin

header of the client CORS request against this list. If there is a match, the request is allowed. To add additional URIs, click the **Add** (plus) icon. To delete a URI from the list, click the **Delete** (minus) icon.

- **Allow Credentials:** If you want the Traffic Router to allow responses to include headers that have the Access-Control-Allow-Credentials header set to true, choose **Yes**.
- **Exposed Headers:** Enter the headers that the client is allowed to access. To add additional headers, click the **Add** (plus) icon. To delete a header from the list, click the **Delete** (minus) icon.
- **Preflight Allowed Methods:** Check the check boxes for the methods that the Traffic Router should allow when a client is accessing the resources of the Delivery Service. This check occurs during a preflight check.
- **Preflight Allowed Headers:** Enter the general headers that the client can use when making the actual request. This check occurs during a preflight request. To add additional headers, click the **Add** (plus) icon. To delete a header from the list, click the **Delete** (minus) icon.
- **Preflight Max Age:** Enter the length of time that a client is allowed to cache the results of a preflight request.

When you are finished editing the Advanced Settings, click **Save**.

Advanced Settings: Geo Limit

Geo Limit:

Geo Limit

CZF + CountryCode(s)

Geo Limit Country Codes

US

Geo Limit Redirect URL

http://rds1.cdn.companyx.com

Anonymous Blocking Enabled?

No

ASN Whitelist

64512, 64515

Save

- **Geo Limit:** The Geo Limit setting enables you to configure the Traffic Router to block client requests based on the coverage zone file (CZF) and the geolocation of the client IP address, depending on the option that you choose. The options for this setting are:
 - **None:** When Geo Limit is set to None, geo blocking is *not* enabled and the Traffic Router will check the Anonymous Blocking configuration to determine whether to allow the request. If Anonymous blocking is not enabled, then ASN blocking is checked to determine whether to allow the request. If ASN blocking is not enabled, the Traffic Router will allow all client requests. None is the default setting for Geo Limit.



Note For more information on Anonymous Blocking, see [Anonymous Blocking, on page 169](#). For more information on ASN Blocking, see [ASN Blocking, on page 173](#).

- **CZF only:** With this setting the Traffic Router will only allow requests from clients whose IP addresses match an entry in the CZF file. If the client IP address does *not* match an entry in the CZF file, then the Traffic Router will reject the client request.



Note The CZF file is discussed in more detail in [Manage Client Routing, on page 155](#). Also an example of the CZF file can be found in [Example CZF File, on page 257](#).

- **CZF + CountryCode(s):** With this setting the Traffic Router will perform the following steps to determine whether a client request is allowed:



Note If you are setting the Geo Limit to "CZF + CountryCode(s)", confirm that the MaxMind geolocation database has been configured in the Media Streamer deployment. For more information on configuring Media Streamer to use the MaxMind database, refer to the "Configure the MaxMind Database for Geolocation" appendix in the *Cisco Media Streamer and Cisco Media Broadcaster Installation and Upgrade Guide*.

1. The Traffic Router checks the CZF file. If the client IP address matches an entry in the CZF file, the Traffic Router will then check the client request against the Anonymous Blocking configuration to determine whether the request is allowed.
2. If the client IP address does *not* match an entry in the CZF file, then the Traffic Router looks up the client IP address in a geolocation database to determine which country the client IP address is in. If the client IP address is in one of the countries listed in the Geo Limit Country Codes field, the Traffic Router will then check the client request against the Anonymous Blocking configuration to determine whether the request is allowed.



Note For more information on Anonymous Blocking, see [Anonymous Blocking, on page 169](#).

3. If the client IP address is not in the CZF file and the client IP address is not in one of the countries listed in the Geo Limit Country Codes field, the Traffic Router checks the NGB whitelist if configured. If the client IP address matches an entry in the NGB whitelist, Anonymous Blocking is checked next to determine whether the request is allowed. If there is no NGB whitelist configured or if the client IP address does not match an entry in the NGB whitelist, the client is redirected to the Geo Limit Redirect URL if it is configured on the Delivery Service. If the Geo Limit Redirect URL is not configured, the client will receive a 503 Service Unavailable message.



Note For more information on the NGB whitelist, see [NGB Whitelist, on page 176](#).

- **Geo Limit Country Codes:** Choose the country codes for the countries you want to permit client requests from. This field is used when Geo Limit is set to **CZF + CountryCode(s)**.
- **Geo Limit Redirect URL:** If the Geo Limit field is set to either "CZF only" or "CZF + CountryCode(s)", you can enter a URL in this field that the Traffic Router will redirect the client to if their request is blocked. If you leave this field blank, the client will receive a "503 Service Unavailable" response.



Note You will only see this field if the Geo Limit is set to either "CZF only" or "CZF + CountryCode(s)".

- **Anonymous Blocking Enabled?:** This setting determines whether this Delivery Service will use the Anonymous Blocking feature to determine whether a client request is allowed. If this value is set to Yes, the Anonymous Blocking profile settings that are configured for the Traffic Router that services this Delivery Service will be used to help determine whether a client request is allowed. For information on how Anonymous Blocking works and how it interacts with the other client blocking methods, see [Understanding CDN Client Blocking Options, on page 23](#). For information on how to configure the Anonymous Blocking settings of the Traffic Router profile, see [Configure Anonymous Blocking, on page 169](#).



Note You cannot enable Anonymous Blocking for DNS Delivery Services.

- **ASN Whitelist:** If the Traffic Router profile is configured to use an ISP database, you can restrict requests to the Delivery Service based on the AS number to which the client IP address belongs. In the ASN Whitelist field, enter a comma separated list of AS numbers from which clients requests should be *allowed*. Client requests from AS numbers not listed in the ASN Whitelist field are blocked.

If the ASN Whitelist field is empty, ASN blocking will *not* block any of the requests.



Note The MaxMind ISP or GeoLite "ASN" database is used to determine which ASN an IP address belongs to. Entries in the ASN White list field only take effect if the Traffic Router profile is configured to use one of these databases. A license is required to use the MaxMind ISP database.

For more information on the assigning an ISP database to the Traffic Router profile, see [ASN Blocking, on page 173](#).

Clone a Delivery Service

To make it easier to add new Delivery Services that are similarly configured, OMD provides the ability to clone (copy) existing Delivery Services to replicate many of their settings. To clone a Delivery Service:

Procedure

-
- Step 1** Choose **Provisioning > Edit CDN**. From the CDN Name drop-down list above the tabs, make sure the correct CDN is selected.
- Step 2** Click the **Delivery Services** tab.
- Step 3** To the far right of the Delivery Service Name drop-down list, click the **Clone** icon to create a Delivery Service.

When you clone a Delivery Service, all settings are copied except for the following settings, which can be edited only until you save the new Delivery Service, but not afterwards:

- Delivery Service Name
 - Display Name
 - Long Description
 - The Delivery Service Subdomain name
-

Steering Delivery Service

A steering Delivery Service is a Delivery Service that is used to “steer” traffic to other Delivery Services. It creates a virtual Delivery Service with an FQDN but no origin mapping. Instead a steering Delivery Service will be assigned target Delivery Services.

The target Delivery Services are assigned a steering type of either order or weight and a value for that type. Whichever steering type you do not assign will have a value of 0.

- Weight values can be 0 or greater.
- Order values can be any integer, including negative values.

The Traffic Router will route any request to a steering Delivery Service FQDN to one of the target Delivery Services, based on the order and weight values.



Note

The target Delivery Services must have a routing type of HTTP and must be part of the same CDN.

Some possible use cases for a steering Delivery Service include:

- Migrating traffic from one Delivery Service to another over time.
- Load balancing between Delivery Services.

- To create an alias for a Delivery Service so that one Delivery Service can have two FQDNs. In this case, there would be only one target Delivery Service for the steering Delivery Service.

Target Delivery Service Selection

The Traffic Router will choose the target Delivery Service by checking the order and weight values in the following order:

1. The Traffic Router will look at target Delivery Services with an order value less than 0, trying them in order from the lowest value to the highest value until an available cache is found. If no available cache is found, the Traffic Router will proceed to Step 2.
2. The Traffic Router will use target Delivery Services with a non-zero weight value. The Traffic Router will distribute requests across the Delivery Services based on the hash of the content URL and the weight of the Delivery Service. This ensures that requests to the same URL are serviced by the same Delivery Service if it has an available cache. Delivery Services with higher weights will receive more requests. If no available cache is found, the Traffic Router will proceed to Step 3.
3. The Traffic Router looks for target Delivery Services with an order value of 0 or greater and tries them in order from the lowest value to the highest value. If no target Delivery Service is available or if there are no available caches for any of the target Delivery Services, the client is redirected to the bypass location listed in the Bypass FQDN field of the Delivery Service. If the Bypass FQDN field is empty, the client will receive a 503 Service Unavailable message.



Note

For Step 1 and Step 3, If two Delivery Services have the same "lowest" order value, the Delivery Service with the alphabetically lowest Delivery Service name is selected.

There are two types of steering Delivery Services:

- **Steering:** For this type of steering Delivery Service, a 302 redirect with a single endpoint URL is returned to the client, based on the order and weight configured on the target Delivery Services.
- **Client Steering:** For this type of steering Delivery Service, a 302 redirect with a list of endpoint URLs based on all of the target Delivery Services is returned to the client. The Location header of the response contains the preferred endpoint URL and the JSON response body contains the URLs listed in order based on their order and weight across all target Delivery Services. If the client request contains the query parameter "trred=false", a 200 response with the JSON response body but no Location header will be sent instead of the 302 redirect with a Location header.

Create a Steering Delivery Service Overview

To create a steering Delivery Service, perform the following steps:

1. Create the target Delivery Services. The target Delivery Services are just regular Delivery Services, but they must have a Routing Type of HTTP. Also make sure that the Origin Server Base URL and the Sub Domain are different for each target Delivery Service. For information on creating the target Delivery Services, see [Add a New DNS or HTTP Delivery Service, on page 92](#).
2. Create the steering Delivery Service. For information on creating a steering Delivery Service, see [Create a Steering Delivery Service, on page 118](#).

3. Assign the target Delivery Services to the steering Delivery Service and assign them Order or Weight values. For information on assigning the target Delivery Services to the steering Delivery Service, see [Assign Target Delivery Service to a Steering Delivery Service, on page 124](#).

Create a Steering Delivery Service

To add a new steering Delivery Service, perform the following steps:



Note

To create a new DNS or HTTP Delivery Service, see [Add a New DNS or HTTP Delivery Service, on page 92](#).

Procedure

- Step 1** Choose **Provisioning > Edit CDN**. From the **CDN Name** drop-down list above the tabs, make sure the correct CDN is selected.
- Step 2** Click the **Delivery Services** tab.
- Step 3** Next to the **Delivery Service Name** drop-down list, click the + icon to create a Delivery Service.
- Step 4** In the **Delivery Service Name**, enter a name for the Delivery Service.
- Step 5** In the **Display Name**, enter a more descriptive display name or this can be the same as the Delivery Service Name.
- Step 6** In the **Long Description** field, enter a description for the Delivery Service.
- Step 7** From the **Use Multi Site Origin Feature** drop-down list, choose Yes to enable MSO for the for this Delivery Service. Otherwise leave it set to the default of No. For more information on the MSO feature, see [Multi Site Origin](#).
- Step 8** The **Delivery Service Profile** drop-down list contains settings for advanced features of the Delivery Service, such as MSO, URL Signing, Regex Remap Parameters, and Edge Geo Blocking. If the Delivery Service you are adding will use any of these advanced features, choose the profile that contains the settings that you would like to use for these features. For more information on Delivery Service profiles, including how to create and edit them, see [Manage Profiles, on page 163](#).
- Step 9** In the Routing Type area, choose one of the following routing types:
 - **STEERING:** For this type of steering Delivery Service, a 302 redirect with a single endpoint URL is returned to the client, based on the order and weight configured on the target Delivery Services. For more information on how the endpoint URL is selected, see [Target Delivery Service Selection, on page 117](#).
 - **CLIENT_STEERING:** For this type of steering Delivery Service, a 302 redirect with a list of endpoint URLs based on all of the target Delivery Services is returned to the client. The Location header of the response contains the preferred endpoint URL and the JSON response body contains the URLs listed in order based on their order and weight across all target Delivery Services. If the client request contains the query parameter "tred=false", a 200 response with the JSON response body but no Location header will be sent instead the 302 redirect with a Location header. For more information on how the ordered of the endpoint URLs is decided, see [Target Delivery Service Selection, on page 117](#).
- Step 10** From the **Protocol** drop-down list, choose one of the following options:
 - **HTTP:** Serves the data to the client using HTTP. Clients trying to use HTTPS will receive an SSL handshake error.

- **HTTPS:** Serves the data to the client using HTTPS. Clients trying to use HTTP will receive a 503: Service Unavailable error.
- **HTTP and HTTPS:** Data will be served to the client using either HTTP or HTTPS, depending on what the client requests.
- **HTTP to HTTPS:** Non-secure (HTTP) clients are redirected with a 302 redirect message to use a secure (HTTPS) connection. Secure (HTTPS) clients will continue with their HTTPS connection.

If you choose any of the options that contain HTTPS, you must also provide the SSL details. Follow these steps to enter the SSL details:

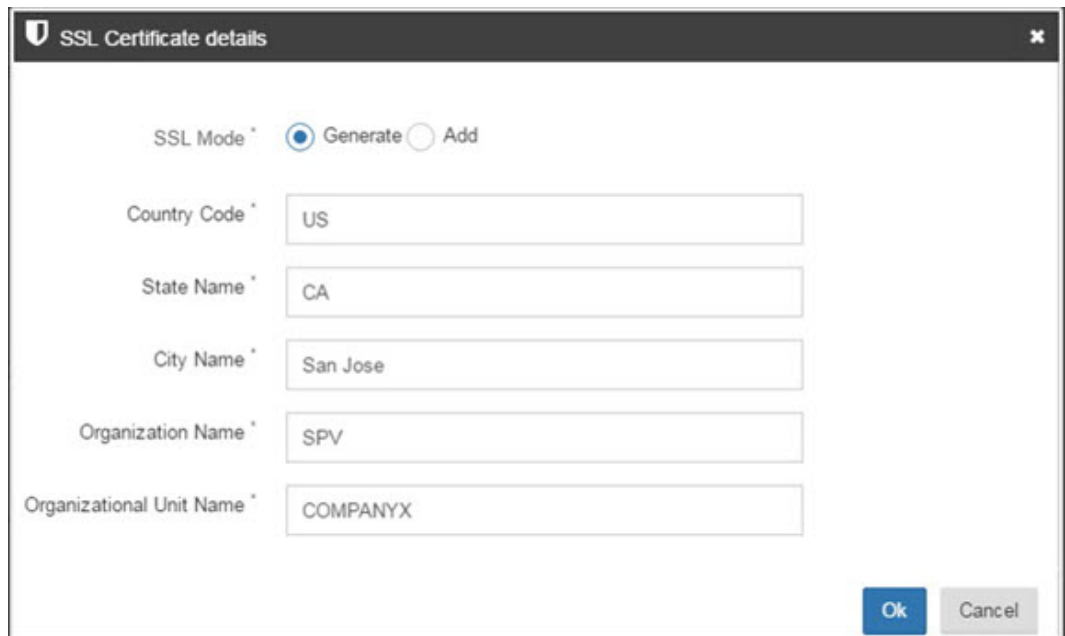
- a) Click the shield icon.

The screenshot shows the 'Provisioning > Edit CDN' interface. At the top, there's a 'CDN Name' dropdown set to 'omd-cdn1' and a 'Delete' button. Below this is a tabbed interface with 'Servers', 'Cache Groups', 'Delivery Services' (selected), 'Client Routing', 'Profiles', and 'Content Invalidation'. The 'Delivery Services' tab contains the following fields:

- Delivery Service Name ***: Text input with value 'sp1-steering'.
- Display Name ***: Text input with value 'sp1-steering'.
- Long Description**: Text area with value 'Steering for SP1'.
- Delivery Service profile ***: Dropdown menu with value 'No Profile'.
- Routing Type ***: Radio buttons for 'DNS', 'HTTP', 'STEERING' (selected), and 'CLIENT_STEERING'.
- Edge cache retrieval from origin**: Dropdown menu with value 'No'.
- Protocol ***: Dropdown menu with value 'HTTPS'. To its right is a shield icon with a tooltip that says 'Click to enter SSL details'.
- Customer**: Empty text input field.
- Active ***: Dropdown menu with value 'No'.

- b) The SSL Certificate Details window appears. For SSL Mode you can generate an SSL certificate or you can add an existing certificate for the Delivery Service.

- **Generate:** If you choose this option, enter all of the required values, indicated by an asterisk (*), in the SSL Certificates Detail window and then click **OK**:



The image shows a dialog box titled "SSL Certificate details" with a close button (X) in the top right corner. Inside the dialog, there are two radio buttons for "SSL Mode": "Generate" (selected) and "Add". Below this, there are five text input fields, each with an asterisk (*) indicating it is a required field. The fields are labeled "Country Code", "State Name", "City Name", "Organization Name", and "Organizational Unit Name". The values entered in these fields are "US", "CA", "San Jose", "SPV", and "COMPANYX" respectively. At the bottom right of the dialog, there are two buttons: "Ok" (blue) and "Cancel" (gray).

SSL Certificate details

SSL Mode * ☒ Generate ☐ Add

Country Code * US

State Name * CA


City Name * San Jose

Organization Name * SPV

Organizational Unit Name * COMPANYX

Ok Cancel

- **Add:** If you choose this option, enter all of the required values, indicated by an asterisk (*), in the SSL Certificates Detail window and then click **OK**:

A dialog box titled "SSL Certificate details" with a close button (X) in the top right corner. It contains three sections, each with a label and a text input field. The first section is "SSL Mode *" with two radio buttons: "Generate" (unselected) and "Add" (selected). The second section is "Private Key *" with a text input field containing the placeholder text "Enter Private Key". The third section is "Certificate Signing Request (CSR)" with a text input field containing the placeholder text "Certificate Signing Request". Below these is a fourth section labeled "Certificate (CRT)*" with a text input field containing the placeholder text "Enter Certificate". At the bottom right of the dialog are two buttons: "Ok" (blue) and "Cancel" (gray).

SSL Certificate details

SSL Mode * ☐ Generate ☒ Add

Private Key *

Certificate Signing Request (CSR)

Certificate (CRT)*

Ok Cancel

Step 11 The **Customer** field is for backwards compatibility only. For OMD Director 3.11 and later, use the Content Provider field instead.

Note If there is a value entered in the **Content Provider** field and the **Customer** field, the Customer field is ignored.

Step 12 From the **Content Provider** drop-down list, choose the Content Provider organization to associate with this Delivery Service. This is an optional field, however, the Content Provider field plays an important role in controlling who can view the OMD Insights analytics information for this Delivery Service.

- **Content Provider Viewers and Content Provider Admins:** OMD Director users that have the role of Content Provider Viewer or Content Provider Admin can only view OMD Insights information for Delivery Services that are assigned a Content Provider that matches the Content Provider organization they are assigned.
- **Reseller Viewers and Reseller Admins:** OMD Director users that have the role of Reseller Viewer or Reseller Admin can only view OMD Insights information for Delivery Services that are assigned a Content Provider that is assigned to the Reseller organization they are assigned.

For example, if user jsmith has the role of Content Provider Admin and has been assigned the Organization of CompanyX, and a Delivery Service named VOD-SJ has been assigned a Content Provider of CompanyY, jsmith *cannot* view any of the OMD Insights information for the VOD-SJ Delivery Service.

Note The **Content Provider** field only affects users that have one of the Content Provider or Reseller roles. It does not limit or restrict users that have been assigned any other role.

Step 13 From the **Active** drop-down list, choose **Yes** to enable the Delivery Service.

Step 14 In the **Sub Domain** field, you can create the sub domain for a Delivery Service in two different ways:

- **Enter only the sub domain:** Enter only the sub domain portion of the FQDN. This subdomain is prepended to the CDN domain and either "tr" (for an HTTP Delivery Service) or "edge" (for DNS Delivery Service) is the hostname used to create the URL for client requests to this Delivery Service. For example, if the CDN domain name is spcdn.company.example and you enter a sub domain of on-site, the client will use a URL for this sub domain of tr.on-site.spcdn.company.example for an HTTP Delivery Service and a URL of edge.on-site.spcdn.company.example for a DNS Delivery Service. The following is an example:

Domains for this Delivery Service:

Sub Domain *	Sample Client URL
on-site	http://tr.on-site.spcdn.company.example

- **Enter a full FQDN:** A custom FQDN enables you to use a different hostname than the default hostname of "tr" for an HTTP Delivery Service or "edge" for the DNS Delivery Service. For this option you enter the full FQDN: hostname, sub domain you want to create, and domain of the CDN. for example srv1.on-site.spcdn.company.example. The following is an example:


Note You can only add one custom FQDN to a Delivery Service.

Figure 1: Delivery Service Custom FQDN Example

Domains for this Delivery Service:

Sub Domain *	Sample Client URL
srv1.on-site.spcdn.company.example	http://srv1.on-site.spcdn.company.example

Step 15 When you are finished entering the information for the Delivery Service, click **Save**. The Processing window will appear. This window displays the progress of the Delivery Service being created, including the status of each action involved.

 Provisioning > Edit CDN

CDN Name :

Servers Cache Groups **Delivery Services** Client Routing Profiles Content Invalidation

Delivery Service Name *

Display Name *

Long Description

Delivery Service profile *

Routing Type * ☐ DNS ☐ HTTP ☒ STEERING ☐ CLIENT_STEERING

Protocol *

Customer

Content Provider


Active *



Domains for this Delivery Service:

Sub Domain * Sample Client URL

Step 16

When the system has successfully created the steering Delivery Service, the Processing status in the title bar will disappear and you will see a green check mark for every action in the Status column, as seen below. Click **OK** to close the window.

Delivery Service create (sp1_steering)  [Show/Hide Config](#)

Action Name	Succeeded	Failed and Reasons	Status
Create delivery_service sp1_steering in Traffic Ops	sp1_steering		
Activate Configuration on Traffic Routers (Snapshot CRConfig)	CDN2		

- Step 17** After the steering Delivery Service is created, you need to assign the target Delivery Services to the steering Delivery Service and assign them Order or Weight values. To accomplish this, perform the steps in [Assign Target Delivery Service to a Steering Delivery Service, on page 124](#).

Assign Target Delivery Service to a Steering Delivery Service

To assign the target Delivery Services to the steering Delivery Service and assign them Order or Weight values, perform the following steps:

Procedure

- Step 1** Choose **Provisioning** > **Edit CDN**. From the CDN Name drop-down list above the tabs, make sure the correct CDN is selected.
- Step 2** Click the **Delivery Services** tab.
- Step 3** From the Delivery Services window, choose the steering Delivery Service to which you are assigning a target Delivery Service from the **Delivery Service Name** drop-down list and click **Manage Steering Assignments**.
- Step 4** From the Delivery Services window make sure the steering Delivery Service is selected from the Delivery Service Name drop-down list and click **Manage Steering Assignments**.
- Step 5** From the Steering Assignment window, choose the first target Delivery Service from the Delivery Service Name drop-down list.
- Step 6** From the **Type** drop-down list, choose **STEERING_WEIGHT** or **STEERING_ORDER** and enter a value in the Value field.

Note For information on how the weight and order values are used to select a target Delivery Service, see [Steering Delivery Service, on page 116](#).

The screenshot shows a window titled "Steering Assignment: sp1_steering". Inside, there are three main fields: "Delivery Service Name" with a dropdown menu showing "sp1-ds1", "Type" with a dropdown menu showing "STEERING_WEIGHT", and "Value" with a text input field containing "1". To the right of the "Value" field is a "Delete" button. Below these fields are three buttons: "Add Target", "Save", and "Cancel".

- Step 7** To add another target Delivery Service, click **Add Target**.
- Step 8** When you are finished adding the target Delivery Services, click **Save**.
- Step 9** After you have finished assigning the target Delivery Services, if you are ready to enable the steering Delivery Service, from the Active drop-down list, choose **Yes** and then click **Save** to save the changes.

Edit an Existing DNS or HTTP Delivery Service

To edit an existing DNS or HTTP Delivery Service, follow these steps:

Procedure

-
- Step 1** Choose **Provisioning > Edit CDN**. From the CDN Name drop-down list above the tabs, make sure the correct CDN is selected.
- Step 2** Click the **Delivery Services** tab.
- Step 3** From the **Delivery Service Name** drop-down list, choose the Delivery Service that you want to edit.
- Step 4** Change the desired settings. For more information on the settings, see [Add a New DNS or HTTP Delivery Service, on page 92](#)
- Step 5** To configure the advanced settings for the Delivery Service, click the **Advanced Settings** link below the **Assign Device Groups to Delivery Service** field. For more information on the advanced settings, see [Delivery Service Advanced Settings, on page 100](#).
- Step 6** When you are finished making changes to the Delivery Service, click **Save**.
- Step 7** When the Delivery Service has been successfully updated, the “Processing” status in the title bar will disappear and you will see a green check mark for every action in the Status column. Click **OK** to close the status window.
-

Edit an Existing Steering Delivery Service

To edit an existing steering Delivery Service, follow these steps:

Procedure

-
- Step 1** Choose **Provisioning > Edit CDN**. From the CDN Name drop-down list above the tabs, make sure the correct CDN is selected.
- Step 2** Click the **Delivery Services** tab.
- Step 3** From the **Delivery Service Name** drop-down list, choose the Delivery Service that you want to edit.
- Step 4** Change the desired settings. For details on the steering delivery service parameters, see [Create a Steering Delivery Service, on page 118](#).
- Step 5** To change which target Delivery Services are assigned to the steering Delivery Service or to change the Type or Value settings of a target Delivery Service, click **Manage Steering Assignments**. In the Steering Assignment window make the necessary changes and click **Save**.
- Step 6** When you are finished making changes to the Delivery Service, click **Save**.
- Step 7** When the system has successfully created the steering Delivery Service, the Processing status in the title bar will disappear and you will see a green check mark for every action in the Status column. Click **OK** to close the window.
-

Delete a Delivery Service


Note

If you are running OMD Director in an HA configuration, you can only make configuration changes from the Primary OMD Director instance. To perform the steps in this section, make sure you are logging into the Primary OMD Director instance.

To delete a Delivery Service, follow these steps:

Procedure

Step 1 Choose **Provisioning > CDN Overview**. From the **CDN Name** drop-down list above the tabs, make sure the correct CDN is selected.

Step 2 In the **Delivery Service** area, click the **Delete** icon (trashcan) for the Delivery Service that you want to delete.

Note If the word “updating” appears instead of the Delete icon (trashcan), this indicates that the Delivery Service has been modified and is in the process of being updated. To delete the Delivery Service, you will need to wait until it has finished updating.

Delivery Service Name	Routing Type	Protocol	Active	Status	Action
vod	http	http	UP	success	
https-dscp	http	https	UP	success	

Step 3 In the Deleting Delivery Service confirmation window that appears, click **Yes** to confirm that you want to delete the Delivery Service.

Note The system is blocked while the changes are made.

Step 4 When the Delivery Service has been successfully deleted, the Processing status in the title bar will disappear and you will see a green check mark for every action in the Status column. Click **OK** to close the status window.

Advanced Delivery Service Features

There are several advanced features you can configure for the Media Streamer Delivery Services. These features include Multi Site Origin (MSO), URL signing, Regex Remap Settings, CDN Routing Name, and Edge Geo Blocking. This section discusses these advanced features and how to configure them.

URL Signing

The Cisco Media Streamer CDN accepts and fulfills requests for video content from client devices in the form of content URLs. Content and service providers, to protect their copyright and fulfill their licensing obligations, use various techniques to protect their content against unauthorized access. URL signing is one such technique. URL signing attaches a keyed hash to the URL, using a secret key shared only between the signer (the portal) and the validating component (Media Streamer cache server).

The URL signing feature of Media Streamer provides the infrastructure to validate content URLs to prevent unauthorized access. URL signing can restrict access by evaluating various parameters embedded in the URL, such as time, unique client identification, and so on. These values can then be validated against the actual client sending the request and the current time at the cache engine serving the request. If any of the validations fail, the request is rejected.

**Note**

Which parameters are evaluated is determined by the Delivery Service profile configuration. If the URL expiry time is used as part of the validation, it is important that the clocks are synchronized on the server that signs the URL and the Edge cache servers that are validating the URL. Network Time Protocol (NTP) should be used on all devices, including the device that is signing the URL.

The URL signing feature of OMD supports both symmetric and asymmetric (public/private) keys for signing and validating content URLs. Symmetric key signing uses the same key to sign and validate the URL. Asymmetric keys always have a key pair, which is made up of a public key and private key: the private key is used for signing and the public key is used for validation.

The following components are used to support the URL Signing feature:

- **Portals:** Portals are responsible for the keys and for signing the URLs. The keys should be provided to the Media Streamer CDN by the portal.

**Note**

The portal is not part of the Media Streamer CDN.

- **Traffic Vault:** Traffic Vaults are the designated keystores (using RIAK) for the Media Streamer CDN. All keys that will be used to validate content URLs using URL signing are stored in the Traffic Vaults. To support the URL Signing feature, you must have Traffic Vault servers installed in the Media Streamer CDN.
- **Traffic Ops:** Traffic Ops acts as the admin user for Traffic Vault. The Portal and Edge cache servers do not interface with the Traffic Vaults directly. To store and retrieve keys from the Traffic Vaults, the Portal and Edge cache servers send requests to Traffic Ops. The Traffic Ops UI can also be used to configure the various settings for the URL signing feature, however the preferred configuration method is OMD Director.
- **Edge Cache Servers:** Edge cache servers validate the URLs and implement the course of action to take if the validation fails, based on the URL signing configuration.

The following is an overview of the steps that need to be performed to configure the URL Signing feature in Media Streamer:

1. Define which values in the URL in addition to the signature should be used to validate the content, which URL schemes are supported, and which URLs to validate. This information is configured in the Delivery

Service profile. For information on how to perform this step, see [Configure Delivery Service Profile for URL Signing, on page 128](#).

2. Enable URL signing on the Delivery Service and identify the keys to be used by the Edge caches that serve the Delivery Service. OMD Director will send the keys to Traffic Ops, and Traffic Ops will store these keys in the Traffic Vaults. For information on how to perform this step, see [Add a New DNS or HTTP Delivery Service, on page 92](#) if you are creating a new Delivery Service or see [Edit an Existing DNS or HTTP Delivery Service, on page 124](#) if you are editing an existing Delivery Service.



Note

The keys that need to be used should be provided by the portal. The keys used for URL signing are shared secret keys between the device that signs the URL (the portal) and the device that validates the URL (the Edge cache). These keys can be either symmetric or asymmetric (public/private) keys.

Configure Delivery Service Profile for URL Signing

The Delivery Service profile is used to configure the URL Signing settings for a Delivery Service. It determines which values in the URL in addition to the signature should be used to validate the content, which schema are supported, and which URLs to validate.

There is only one Delivery Service profile that is automatically created as part of the Media Streamer installation. If different Delivery Services will use different URL Signing settings, you will need to create additional Delivery Service profiles. Also the Delivery Service profile contains settings for additional features, such as MSO, Regex Remap Parameters, and Edge Blocking. If two Delivery Services will use the same URL Signing settings and the same settings for all of these additional features, they can use the same Delivery Service profile.

Perform the following steps to configure a Delivery Service profile for URL Signing:

Procedure

- Step 1** Choose **Provisioning > Edit CDN**.
- Step 2** Make sure the correct CDN is selected in the CDN Name drop-down list and click the **Delivery Services** tab.
- Step 3** From the Delivery Service Name drop-down list, choose the Delivery Service to configure for URL Signing.
- Step 4** In the Delivery Service Profile drop-down list, determine whether a profile is currently assigned to the Delivery Service. If no profile is currently assigned to the Delivery Service, go to Step 5 to create a new Delivery Service profile. If there is a profile currently assigned to the Delivery Service, perform the following steps to configure the URL Signing settings of the profile:
 - a) Choose **Provisioning > Edit CDN**.
 - b) Make sure the correct CDN is selected in the CDN Name drop-down list and click the **Profiles** tab.
 - c) From the Profile Details window, click **Edit** in the Action column for the profile that is assigned to the Delivery Service for which you are configuring URL Signing.
 - d) Click the **URL Signing** tab.
 - e) The Delivery Service Settings window will display the current URL Signing settings assigned to the profile. These settings will apply to any Delivery Service to which you assign this profile. Make changes to the fields as needed:
 - **Validate Client IP:** Check this check box to include the client IP address as part of the URL validation.

- **Validate Expiry Time:** Check this check box to include the Expiry time as part of the URL validation.
- **Validate Domain Name:** Check this check box to include the Domain name as part of the URL validation.
- **Valid Schema:** From this drop-down list, choose which schemes are valid for URLs that are being validated to use. The options are All, HTTP, and HTTPS.
- **Regular Expression:** In this field, enter a regular expression to identify which URLs to validate. For example, the following regular expression could be an example that would cause the Edge cache servers to validate the HLS master manifest: `\(?:!(S|s)tream)[^v]*\.(m3u8|mpd|isml)|manifest$`.
- **ABR Fragment Regular Expression:** This field is used with the Session Tracking Security Enablement feature. In this field, enter a regular expression to identify which subsequent requests *after* the manifest URL for which the Edge cache server should validate the cookie signature. For example, the following regular expression would cause the Edge cache servers to validate the cookie signature on all ABR fragments: `(.*)\\.(ts|mp4|mp4v|mp4a)`.

If the Session Tracking Security Enablement feature is enabled, for any subsequent requests *after* the manifest URL that match the value in the ABR Fragment Regular Expression field, the Edge cache server must validate the cookie signature in that request.

If the cookie signature validation fails, the request is rejected and a 403 HTTP response is sent back to the client. Any subsequent client requests that come in that do *not* match the value in the ABR Fragment Regular Expression field or the URL Signing Regular Expression field of the Delivery Service profile are passed through without further URL checking.

Note If Session Tracking Security Enablement is enabled and the ABR Fragment Regular Expression field is left blank, a default regular expression of `(.*)\\.(ts|mp4|mp4v|mp4a)` is used.

The screenshot shows the 'Provisioning > Edit CDN' interface. At the top, there's a 'CDN Name' dropdown set to 'omd-cdn1' and a 'Delete' button. Below this are tabs for 'Servers', 'Cache Groups', 'Delivery Services', 'Client Routing', 'Profiles', and 'Content Invalidation'. The 'Profiles' tab is active, showing 'Delivery Service Settings - tt1'. Under this, there are sub-tabs: 'Multi Site Origin', 'URL Signing' (which is highlighted with a red box), 'Regex Remap Settings', and 'CDN Routing'. In the 'URL Signing' section, there are several settings: 'Validate Client IP' (checked), 'Validate Expire Time' (checked), 'Validate Domain Name' (checked), 'Valid Schema' (set to 'All'), 'Regular Expression' (with a hint for HLS master manifest), and 'ABR Fragment Regular Expression'. At the bottom right are 'Cancel' and 'Save' buttons.

- f) When you are finished making changes, click **Save**.
- g) When the profile has been successfully updated, the Processing status in the title bar will disappear and you will see a green check mark for every action in the Status column. Click **OK** to close the status window.

Step 5

If no Delivery Service profile was assigned to the Delivery Service, perform the following steps to create a new Delivery Service profile to use:

- a) Choose **Provisioning > General Settings** and click the **Profiles** tab.
- b) In the Action column header, click the **Add Profile** (plus) icon.
- c) In the Create/Copy Profile window that appears, enter the following to create a new Delivery Service profile:
 - **Mode:** Choose **Create** or **Copy**:
 - **Create:** To create a new profile based on the default settings for the profile, choose **Create**. You can edit these values after you create the profile.
 - **Copy:** To create a new profile based on the settings of an existing profile, choose **Copy**. You can edit these values after you create the profile.
 - **Name:** Enter a unique descriptive name for the profile. It is helpful to begin the name with **DS** to easily identify it as a Delivery Service profile.
 - **Description:** Enter a description for the new profile. This description appears when you view the list of available profiles on the Provisioning > Edit CDN > Profiles tab so enter a description that helps you identify this profile.
 - **CDN:** In the CDN drop-down list, make sure the CDN for the Delivery Service you are configuring the URL Settings for is selected.

Note A profile can belong to only one CDN.

- **Type:** If the mode you chose was Create, from the Type drop-down list choose **delivery_service**. If the mode you chose was Copy, you will not see this field.
 - **Copy From:** If the mode you chose was Copy, from the Copy From drop-down list choose the Delivery Service profile from which to create this profile. When you create a profile by copying it from another profile, the new profile will inherit the parameters and settings of that profile. After you finish creating the new profile, you can edit these settings. If the mode you chose was Create, you will not see this field.
- d) Click **Save** to create and save the new profile.
 - e) To view, and if necessary, change the settings of the new profile, choose **Provisioning > Edit CDN**.
 - f) Make sure the correct CDN is selected in the CDN Name drop-down list and click the **Profiles** tab.
 - g) From the Profile Details window, click **Edit** in the Action column for the new profile that you created.
 - h) Click the **URL Signing** tab.
 - i) The Delivery Service Settings window will display the current URL Signing settings assigned to the profile. These settings will apply to any Delivery Service to which you assign this profile. Make changes to the fields as needed:
 - **Validate Client IP:** Check this check box to include the client IP address as part of the URL validation.
 - **Validate Expiry Time:** Check this check box to include the Expiry time as part of the URL validation.
 - **Validate Domain Name:** Check this check box to include the Domain name as part of the URL validation.
 - **Valid Schema:** From this drop-down list, choose which schemes are valid for URLs that are being validated to use. The options are All, HTTP, and HTTPS.
 - **Regular Expression:** In this field, enter a regular expression to identify which URLs to validate. For example, the following regular expression could be an example that would cause the Edge cache servers to validate the HLS master manifest: `V(?!(S|s)treame)[^\\]*\\.(m3u8|mpd|isml)|manifest$`.
 - **ABR Fragment Regular Expression:** This field is used with the Session Tracking Security Enablement feature. In this field, enter a regular expression to identify which subsequent requests *after* the manifest URL for which the Edge cache server should validate the cookie signature. For example, the following regular expression would cause the Edge cache servers to validate the cookie signature on all ABR fragments: `(.*)\\.(ts|mp4|mp4v|mp4a)`.

If the Session Tracking Security Enablement feature is enabled, for any subsequent requests *after* the manifest URL that match the value in the ABR Fragment Regular Expression field, the Edge cache server must validate the cookie signature in that request.

If the cookie signature validation fails, the request is rejected and a 403 HTTP response is sent back to the client. Any subsequent client requests that come in that do *not* match the value in the ABR Fragment Regular Expression field or the URL Signing Regular Expression field of the Delivery Service profile are passed through without further URL checking.

Note If Session Tracking Security Enablement is enabled and the ABR Fragment Regular Expression field is left blank, a default regular expression of `(.*)\\.(ts|mp4|mp4v|mp4a)` is used.

- j) When you are finished making changes, click **Save**.
- k) When the profile has been successfully updated, the Processing status in the title bar will disappear and you will see a green check mark for every action in the Status column. Click **OK** to close the status window.

- Step 6** If you created a new Delivery Service profile to use for URL Signing, perform the following steps to assign the new profile to the Delivery Service:
- a) Choose **Provisioning > Edit CDN**.
 - b) Make sure the correct CDN is selected in the **CDN Name** drop-down list and click the **Delivery Services** tab.
 - c) From the **Delivery Service Name** drop-down list, choose the Delivery Service for which you are configuring Edge Geo Blocking.
 - d) From the **Delivery Service Profile** drop-down list, choose the new profile that you created.
 - e) Click **Save** to save the changes.

Session Tracking Security Enablement

A limitation of URL signing is that while it is easy to enforce the signature in the initial request, which is typically a request for the top level Manifest through a portal, it is difficult to enforce the signature in the subsequent request URLs (such as segment requests). This is because the subsequent request URLs are embedded in the manifest itself, which are generated upstream by a packager/transcoder. Because of this behavior, it is typical to configure the URL signing feature to only check the signature on the initial manifest request. (This is accomplished by creating a regular expression in the URL Signing configuration, such as `@pparam=V(?:S|s)treame)[^\\]*\\.(m3u8|mpd|ismv|)|manifest$.`)

Unfortunately, only validating the requests for the initial manifest request makes it possible for a client to directly request segments from the CDN without first requesting the top level manifest, thereby completely avoiding and bypassing URL signature validation. The Session Tracking Security Enablement, which is an enhancement to the Session Tracking feature, solves this problem.

With Session Tracking Security Enablement enabled, after a URL with a signature is received and validated by the Edge Cache Server, a client-specific cookie is generated and sent back to the client using the 'SET-COOKIE' header. Similar to the URL signature, this cookie contains a cookie signature generated by SHA1 HMAC, using an HMAC-key formed by concatenating the client IP address and a secret key. Using the client IP address in the key makes sure the cookie is specific only to the requesting client.

The client is then expected to send this cookie in subsequent HTTP requests. Upon receiving subsequent client requests (that do not contain a URL signature), the Edge cache server validates the cookie signature for any requests coming in that match the value configured in the ABR Fragment Regular Expression field of the Delivery Service profile. If the validation fails, the request is rejected and a 403 HTTP response is sent back to the client. Any subsequent client requests that come in that do not match the value in the ABR Fragment Regular Expression field or the URL Signing Regular Expression field of the Delivery Service profile are passed through without further URL checking.

The Session Tracking Security Enablement feature is an enhancement to the Session Tracking feature. To use the Session Tracking Security Enablement feature you must have both the URL Signing feature configured and enabled and you must have the Session Tracking feature enabled. For information on configuring Session Tracking Security Enablement, see the [Advanced Settings: Content Preposition](#) section. Information on configuring the ABR Fragment Regular Expression field and the URL Signing Regular Expression settings of the Delivery Service profile are covered in the [Configure Delivery Service Profile for URL Signing](#) section.

Multi Site Origin

Without the Multi Site Origin (MSO) feature configured, Mid cache servers are not aware of any redundancy of the Origin Servers. The MSO feature enables the Mid cache servers to be aware that there are multiple Origin Servers available to serve the Delivery Service, which enables them to provide Origin Server failover and load-distribution features.

You can configure the MSO feature to support failure for both Simple retries (on 4xx HTTP responses) and Dead Server retries (on 5xx HTTP responses):

- **Simple retry:** When the MSO feature is configured, when a 404 response is received, the Mid cache will retry Origin Servers until a successful response is obtained. Each Origin Server is tried multiple times in succession. If all Origin Servers have been tried without success, the client receives a 404 response. The order in which the Origin Servers are tried and how many times they are retried is based on the configuration.
- **Dead server retry:** For dead server retry, a Mid cache checks the response from an Origin Server. If the response received matches one of the Unavailable Server Retry Responses that is configured, the Mid cache will retry Origin Servers until a successful response is obtained from one of the servers, or until all Origin Servers have been tried without success. Each Origin Servers is tried multiple times in succession. An Origin Server that returns a 503 message will be marked unavailable for parent selection for a period of 5 minutes by default. The order in which the Origin Servers are tried and how many times they are retried is based on the configuration.

In OMD Director, you can choose between four different algorithms to provide load distribution and failover for the MSO feature:

- **strict:** With this setting, Traffic Servers will serve requests in a round robin fashion, spreading requests across multiple parents simultaneously based on the order of requests. This method is also referred to as Strict Round-Robin.
- **true:** This setting is essentially the same as strict, but it also ensures that requests from the same IP address always go to the same Origin Server, if that server is available. This method is also referred to as IP-based Round-Robin.
- **consistent_hash:** Spreads requests across multiple parents simultaneously based on the hash of the content URL. The weight that is assigned to the Origin Servers, through the Origin Server profile that is applied to the server, determines the percentage of requests that go to each server in the origin group. For example, if each server in the group is assigned the same weight, the requests would be spread equally among the servers in the group.
- **false:** With this setting, round robin selection does not occur. Instead, only the primary Origin Server is used as long as it is available. If the primary server is no longer available, then a backup Origin Server is used. The order in which Origin Servers are used is based on the configuration. This method is also referred to as Primary/Backup failover.

When the Multi Site Origin algorithm is set to “False” there are two ways to configure which Origin Servers the Mid cache servers use and which Origin Servers are used as primary or backup:

- **Mid Cache Group configuration:** You can assign specific Origin Servers to Mid cache groups by assigning the Origin Servers to origin groups and then assigning those groups as the Parent and Secondary Parent cache groups of the Mid cache group. There are two approaches you can use with this method:
 - Create one origin group and assign both the primary and backup server to the same group. Assign this origin group as the Parent cache group of the Mid cache group. The server with the lowest rank

in this origin group will be the primary and the other server will be the backup. The rank of a server is assigned through the profile that is applied to the server. With this method you do not need to assign a Secondary Parent cache group. If you do assign an origin group as the Secondary Parent cache group, the Origin Servers in this group will be used as backups after the servers in the Parent cache group.

- Create two origin groups and assign the primary Origin Server to one group and the backup Origin Server to the other group. Then assign the group with the primary Origin Server as the Mid cache group's Parent cache group and assign the group with the backup Origin Server as the Secondary Parent cache group. In this configuration, the rank of each server is not considered. With this configuration, you can configure different Mid cache groups to use different primary and different backup Origin Servers. This enables you to provide load distribution, which reduces the traffic volume to a single Origin Server, and still provides failover. For example, you could have Origin Server X be the Primary for Mid cache group X and Origin Server Y be the Primary for Mid cache group Y. Then you can assign Origin Server X to be the backup for Mid cache group Y and Origin Server Y to be the backup for Origin Server X.
- **Delivery Service configuration:** If a Mid cache group that supports the Delivery Service does not have a Parent Cache Group assigned, then only the rank of the Origin Server is used to determine which Origin Server is the Primary. All Mid cache servers that support the Delivery Service that are not assigned a Parent Cache Group will use the same primary Origin Server and same backup Origin Server. The rank for an Origin Server is configured in its profile. The Origin Server with the lowest rank is considered the primary. If two Origin Servers have the same lowest rank, then either Origin Server could end up being used as the primary. With this method, Origin Servers are assigned to a Delivery Service.

Multi Site Origin Configuration Overview

The following is an overview of the steps you must perform to configure Multi Site Origin for a Delivery Service:



Note

For details on configuring these steps, see [Configure Multi Site Origin, on page 135](#).

1. **Create Origin Cache Groups:** If you are assigning the Origin Servers to Mid Cache groups you need to create origin groups.
2. **Create Origin Server Profiles:** The Origin Server profile is used to assign ranks to the Origin Server. The rank is used to help determine the order in which Origin Servers should be tried.
3. **Add Origin Servers to OMD Director:** Typically the Origin Servers will not automatically be registered with the CDN because they are not salt minions, so you will need to add any Origin Servers that you are using for MSO to the CDN. When you add an Origin Server to the CDN you will also assign it an Origin Profile to use and you will assign it to an Origin Server cache group.
4. **Assign Origin Server Cache Groups to Mid Cache Groups:** When the Multi Site Origin algorithm is set to "False", for different Mid cache groups that serve the same Delivery Service to use different Primary and Backup Origin Servers, you need to assign Origin Server Cache Groups to the Parent Cache Group and optionally Secondary Cache Group settings of the Mid Cache Group.
5. **Create or Edit a Delivery Service Profile:** The Delivery Service profile is used to configure the MSO settings for a Delivery Service, such as algorithm and whether to provide failover for both Simple retries

(on 4xx HTTP responses) and Dead Server retries (on 5xx HTTP responses). If different Delivery Services will use different MSO settings, you will need to have more than one Delivery Service profile.

6. **Configure the Delivery Service for MSO:** To configure the Delivery Service to use MSO, you need to enable MSO, assign a profile to the Delivery Service, and optionally assign Origin Servers to the Delivery Service.

Configure Multi Site Origin

The following sections provide the details to configure MSO for a Delivery Service in your CDN.



Note

If you are running OMD Director in an HA configuration, you can only make configuration changes from the Primary OMD Director instance. To perform the steps in this section, make sure you are logging into the Primary OMD Director instance.

View Origin Server Cache Groups

To be able to assign Origin Servers directly to Mid Cache Groups you must assign the Origin Servers to an Origin cache group. Perform the following steps to view the existing Origin cache groups to determine if an Origin cache group already exists that can be used:

Procedure

- Step 1** Choose **Provisioning > CDN Overview**. From the **CDN Name** drop-down list above the tabs, make sure the correct CDN is selected.

The **Origin Cache Groups** area lists the Origin cache groups that are configured in the CDN, including the following information:

- Group name
- Latitude and longitude of the Origin cache group
- Status of the group

You can sort the list by any of these columns by clicking the arrow icons in the column header. From this section you can also view additional details about the Origin cache groups or delete an Origin cache group.

Origin Cache Groups

10

Group Name	Latitude	Longitude	Status	Action
orggrp	-12	-52	success	<div><div></div><div></div></div>

Showing 1 to 1 of 1 entries

Previous

1

Next

- Step 2** To view the details of a Origin cache group, click the **Information** icon (i) for the Origin cache group that you want to view.

Origin Cache Groups				
<input type="text"/> 10				
Group Name	Latitude	Longitude	Status	Action
orggrp	-12	-52	success	
Showing 1 to 1 of 1 entries				
Previous 1 Next				

Step 3 In the window that appears, you can view the settings for the Origin cache group. The following is an example:

Origin Cache Group - orggrp ✕

Group Name	orggrp
Short Name	org-grp
Type	ORG_LOC
Latitude	-12
Longitude	-52
Assigned Servers	org-server, org-server2
Parent Cache Group	CtrlPlaneGroup
Secondary Parent Cache Group	NO_PARENT

Close

Step 4 When you are done viewing the details of the Origin cache group, click **Close**.

Step 5 If you need to create a new Origin cache group, see [Create Origin Server Cache Groups, on page 136](#).

Create Origin Server Cache Groups

To assign Origin Servers directly to Mid Cache Groups you must assign the Origin Servers to an Origin cache group. Perform the following steps to create cache groups for Origin Servers:

Procedure

- Step 1** Choose **Provisioning** > **Edit CDN**. From the CDN Name drop-down list above the tabs, make sure the correct CDN is selected.
- Step 2** Click the **Cache Groups** tab.
- Step 3** Next to the **Cache Group Name** drop-down list, click the + icon to create a new cache group.
- Step 4** In the New Cache Group window that appears, enter the following information for the new group:

- In the **Name** field, enter a descriptive name for the group. This is the name that appears in the OMD Director drop-down menus. To more easily identify this group as a group for Origin Servers, it is recommended that you being the name with ORG.
- In the **Short Name** field, enter an additional descriptive name. This name is not currently used in OMD Director, but it is a required field. The Short Name can be the same as the Name.
- The **Geo Magnetic Latitude** and **Geo Magnetic Longitude** values are not used for Origin Servers, but are required fields, you can enter any valid value in these fields.
- From the **Type** drop-down list, choose **ORG_LOC** to configure this as an Origin Server group. You can only assign Origin Servers to groups of this type.

Note The **Parent Cache Group** and **Secondary Parent Cache Group** fields are set to NO_PARENT and are not editable.

Step 5 Click **Add** to save and add the group.

Step 6 A window showing that the cache group was successfully created appears. Click **OK** to close this window.

Step 7 Repeat Step 1 to Step 6 to create additional Origin cache groups as needed. Refer to [Multi Site Origin, on page 133](#) for more information on how many groups to create based on your desired results.

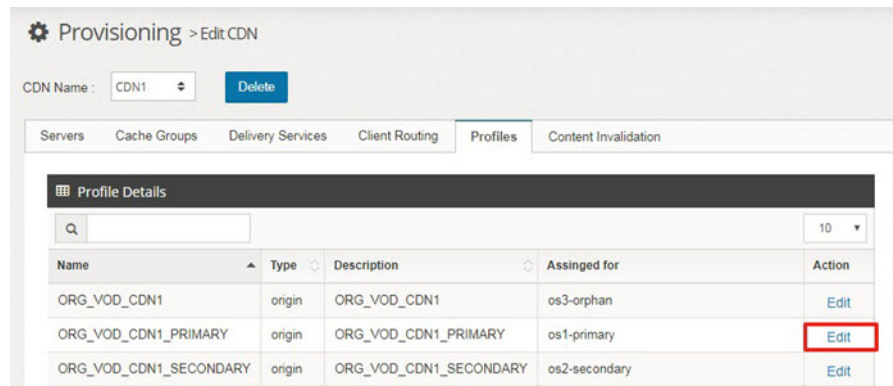
View Origin Server Profiles

The Origin Server profile is used to assign ranks to the Origin Server. The rank is used to determine the order in which Origin Servers should be tried for any Mid cache servers that do not have a Parent Cache Group assigned or if there is more than one Origin Server in the Parent Cache Group. In this case, the Origin Server with the lowest rank is considered the primary. If two Origin Servers have the same lowest rank, then either Origin Server could end up being used as the primary.

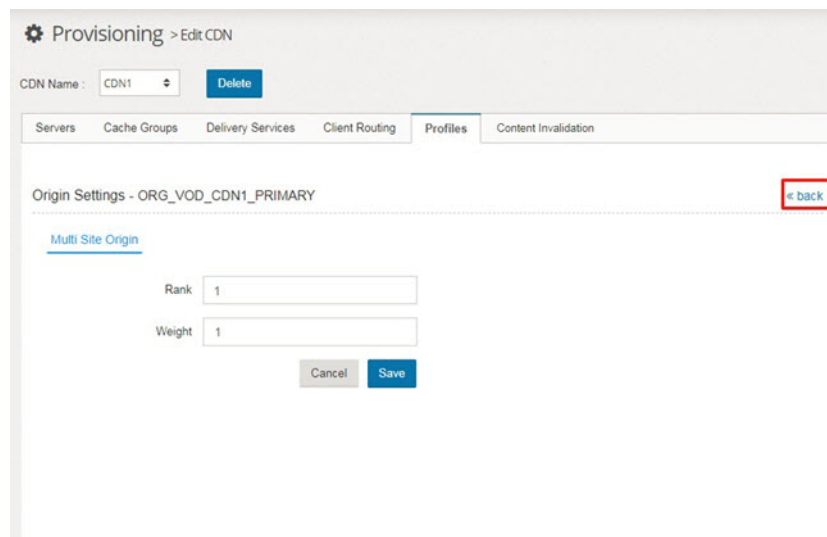
There is only one Origin Server profile that is automatically created as part of the Media Streamer installation. To assign different ranks to different Origin Servers you need to create additional Origin Server profiles.

Procedure

- Step 1** Choose **Provisioning > Edit CDN**.
- Step 2** Make sure the correct CDN is selected in the CDN Name drop-down list and click the **Profiles** tab.
- Step 3** From the Profile Details window, click **Edit** in the Action column for the profile you want to view. The Origin Server profiles should start with ORG_.



- Step 4** The Origin Settings window will display the rank and weight assigned to that profile:
- **rank:** When the MSO algorithm is set to False, this parameter is used to help determine which server should be the primary and which server should be the backup.
 - **weight:** When the MSO algorithm is set to consistent hash, this parameter is used to determine what percentage of requests goes to each server in the origin group.
- Step 5** When you are finished viewing the settings for the profile, click **back** to go back to the list of profiles.



What to do next

If an Origin Server profile does not exist with the rank setting that you need, perform the steps in [Create Origin Server Profile, on page 139](#).

Create Origin Server Profile

If an Origin Server profile does not exist with the rank setting that you need, perform the following steps to create a new Origin Server profile:

Procedure

-
- Step 1** Go to **Provisioning > General Settings** and click the **Profiles** tab.
- Step 2** In the Action column header, click the **Add Profile** (plus) icon.
- Step 3** In the Create/Copy Profile window that appears, enter the following information to create a new Origin Server profile:
- **Mode:** Choose **Create** or **Copy**:
 - **Create:** To create a new profile based on the default settings for the profile, choose **Create**. You can edit these values after you create the profile.
 - **Copy:** To create a new profile based on the settings of an existing profile, choose **Copy**. You can edit these values after you create the profile.
 - **Name:** Enter a unique descriptive name for the profile. It is helpful to begin the name with **OS** to easily identify it as an Origin Server profile.
 - **Description:** Enter a description for the new profile. This description appears when you view the list of available profiles on the Provisioning > Edit CDN > Profiles tab so enter a description that helps you identify this profile.
 - **CDN:** In the CDN drop-down list, make sure the CDN for the Delivery Service you are configuring MSO for is selected.
- Note** A profile can belong to only one CDN.
- **Type:** If the mode you chose was Create, from the Type drop-down list choose **origin**. If the mode you chose was Copy, you will not see this field.
 - **Copy From:** If the mode you chose was Copy, from the Copy From drop-down list choose the Origin Server profile from which you would like to create this profile. When you create a profile by copying it from another profile, the new profile will inherit the parameters and settings of that profile. After you finish creating the new profile, you can edit these settings. If the mode you chose was Create, you will not see this field.
- Step 4** Click **Save** to create and save the new profile.

Create/Copy Profile

Mode * ☐ Create ☒ Copy

Name *

Description *

CDN

Copy From

- Step 5** A status window appears that shows the profile was created successfully. Click **Close** to close this window.
- Step 6** To view, and if necessary, change the settings of the new profile, go to **Provisioning > Edit CDN**.
- Step 7** Make sure the correct CDN is selected in the CDN Name drop-down list and click the **Profiles** tab.
- Step 8** From the Profile Details window, click **Edit** in the Action column for the new profile that you created.
- Step 9** In the Origin Settings window, enter the desired rank and weight for this profile, based on the following information:
- **rank:** When the MSO algorithm is set to False, this parameter is used to help determine which server should be the primary and which server should be the backup. The primary and backup Origin Servers are determined as follows:
 - If the Parent cache group for a Mid cache group contains only one Origin Server, then that Origin Server is the Primary. The Origin Server that is assigned to the Secondary Parent cache group is the backup.
 - If the Parent cache group of a Mid cache group has more than one Origin Server, the Origin Server in that group with the lowest rank is considered the primary and the Origin Server with the second lowest rank is the first backup. If two Origin Servers have the same lowest rank, then either Origin Server could end up being used as the primary, so make sure to configure different ranks if you want to control which server is primary.
 - If no Parent cache group is assigned to the Mid cache group, then the Origin Server with the lowest rank that is assigned to the Delivery Service, either specifically in the Assign Cache Servers list or as part of a device group, is the primary. The Origin Server with the second lowest rank is the first backup. If two Origin Servers have the same lowest rank, then either Origin Server could end up being used as the primary, so make sure to configure different ranks if you want to control which server is primary.
 - **weight:** When the MSO algorithm is set to consistent hash, this parameter is used to determine what percentage of requests go to each server in the origin group. If two servers have equal weights, they will receive an equal percentage of the requests.

- Step 10** When you are finished making changes, click **Save**.

The screenshot shows the 'Provisioning > Edit CDN' interface. At the top, there's a 'CDN Name' dropdown set to 'CDN1' and a 'Delete' button. Below this are tabs for 'Servers', 'Cache Groups', 'Delivery Services', 'Client Routing', 'Profiles', and 'Content Invalidation'. The 'Profiles' tab is active, showing 'Origin Settings - ORG_VOD_CDN1_PRIMARY'. Under 'Multi Site Origin', there are input fields for 'Rank' (set to 1) and 'Weight' (set to 1). At the bottom right, there are 'Cancel' and 'Save' buttons, with the 'Save' button highlighted by a red box.

- Step 11** When the profile has been successfully updated, the “Processing” status in the title bar will disappear and you will see a green check mark for every action in the Status column. Click **OK** to close the status window.
- Step 12** If you need to create additional Origin Server profiles, repeat Steps 1–11.

Add Origin Servers to the CDN for MSO

Typically the Origin Servers will not automatically be registered with the CDN because they are not salt minions, so you will need to add any Origin Servers that you are using for MSO to the CDN. When you add an Origin Server to the CDN you will also assign it an Origin Server Profile to use and you will assign it to an Origin Server cache group. For information on adding Origin Servers to the CDN, see the [Add an Origin Server to the CDN](#) section.

Assign Origin Cache Groups to Mid Cache Groups

When the Multi Site Origin algorithm is set to “False”, if you want different Mid cache groups that serve the same Delivery Service to use different Primary and Backup Origin Servers, you need to assign Origin Server Cache Groups to the Parent Cache Group, and optionally the Secondary Cache Group of the Mid Cache Group.

Perform the following steps to assign an Origin Server cache group to an existing Mid cache group:

Procedure

- Step 1** Choose **Provisioning > Edit CDN**. From the CDN Name drop-down list above the tabs, make sure the correct CDN is selected.
- Step 2** Click the **Cache Groups** tab.
- Step 3** From the Cache Group Name list, choose the Mid cache group to which you want to assign the Origin cache groups.
- Step 4** From the Parent Cache Group drop-down list, choose the Origin cache group that contains the Origin Servers that you want to assign.
- Step 5** If you have created a different Origin cache group to use for your backup Origin Servers, from the Secondary Parent Cache Group drop-down list, choose that Origin cache group.

- Step 6** When you have finished making your changes, click **Save**. The Processing window will appear. This window displays the progress of the Cache Group being updated, including the status of each action involved.
- Step 7** When the Cache Group has been successfully updated, the “Processing” status in the title bar will disappear and you will see a green check mark for every action in the Status column. Click **OK** to close the status window.

Configure Delivery Service Profile for MSO

The Delivery Service profile is used to configure the MSO settings for a Delivery Service, such as algorithm and whether to provide failover for both Simple retries (on 4xx HTTP responses) and Dead Server retries (on 5xx HTTP responses). There is only one Delivery Service profile that is automatically created as part of the Media Streamer installation. If different Delivery Services will use different MSO settings, you will need to create additional Delivery Service profiles. The Delivery Service profile also contains settings for additional features, such as URL Signing, Edge Blocking, and Regex Remap Parameters. If two Delivery Services will use the same MSO settings and the same settings for these additional features, they can use the same Delivery Service profile.

Perform the following steps to configure a Delivery Service profile for MSO:

Procedure

- Step 1** Choose **Provisioning > Edit CDN**.
- Step 2** Make sure the correct CDN is selected in the CDN Name drop-down list and click the **Delivery Services** tab.
- Step 3** From the Delivery Service Name drop-down list, choose the Delivery Service to configure for MSO.
- Step 4** In the Delivery Service Profile drop-down list, determine whether a profile is currently assigned to the Delivery Service. If no profile is currently assigned to the Delivery Service, go to Step 5 to create a new Delivery Service profile. If there is a profile currently assigned to the Delivery Service, perform the following steps to configure the MSO settings of the profile:
- Choose **Provisioning > Edit CDN**.
 - Make sure the correct CDN is selected in the CDN Name drop-down list and click the **Profiles** tab.
 - From the Profile Details window, click **Edit** in the Action column for the profile that is assigned to the Delivery Service for which you are configuring MSO.
 - Click the **MSO** tab.
 - The Delivery Service Settings window will display the current MSO settings assigned to the profile. These settings will apply to any Delivery Service to which you assign this profile. Make changes to the fields as needed:
 - **Algorithm:** Choose one of the following algorithms based on the MSO failover and load-balancing approach you want the Delivery Services that use this profile to use:
 - **strict:** With this setting, Traffic Servers will serve requests in a round robin fashion, spreading requests across multiple parents simultaneously based on the order of requests. This method is also referred to as Strict Round-Robin.
 - **true:** This setting is essentially the same as strict, but it also ensures that requests from the same IP address always go to the same Origin Server, if that server is available. This method is also referred to as IP-based Round-Robin.

- **false:** With this setting, round robin selection does not occur. Instead, only the primary Origin Server is used as long as it is available. If the primary server is no longer available, then a backup Origin Server is used. The order in which Origin Servers are used is based on the configuration. This method is also referred to as Primary/Backup failover.
- **consistent_hash:** Spreads requests across multiple parents simultaneously based on the hash of the content URL.
- **Parent Retry:** This setting determines whether failover of Origin Servers should be provided for 404 HTTP responses (simple retry) or 5xx HTTP responses (dead server retry):
 - **Both:** Failover is provided for both 404 HTTP responses and any 5xx HTTP responses that are configured in the “Unavailable Server Retry Response” list in the Delivery Service profile.
 - **Simple Retry:** Failover is only provided for 404 HTTP responses.
 - **Unavailable Server Retry:** Failover is only provided for the 5xx HTTP response codes that are configured in the “Unavailable Server Retry Response” list in the Delivery Service profile.
- **Unavailable Server Retry Response:** If you selected either “both” or “Unavailable Server Retry” for Parent Retry, then in this field enter the 5xx HTTP responses that should trigger failover. Only enter one value per field. To add extra response codes, click the + icon after the last response code listed and a new field will appear. To remove one of the codes, click the – icon.
- **Max Simple Retries:** Determines the maximum number of times an Origin Server should be retried for any 404 HTTP responses.
- **Max Unavailable Server Retries:** Determines the maximum number of times an Origin Server should be retried for any of the 5xx HTTP responses configured in the Unavailable Server Retry Response setting.

Figure 2: Edit a Delivery Service Profile

Provisioning > Edit CDN

CDN Name :

Servers Cache Groups Delivery Services Client Routing **Profiles** Content Invalidation

Delivery Service Settings - ds_profile

Multi Site Origin URL Signing Regex Remap Settings

Algorithm

Parent Retry

Unavailable Server Retry Response

Max Simple Retries

Max Unavailable Server Retries

- f) When you are finished making changes, click **Save**.
- g) When the profile has been successfully updated, the “Processing” status in the title bar will disappear and you will see a green check mark for every action in the Status column. Click **OK** to close the status window.

Step 5

If there was no Delivery Service profile assigned to the Delivery Service, perform the following steps to create a new Delivery Service profile to use:

- a) Choose **Provisioning > General Settings** and click the **Profiles** tab.
- b) In the Action column header, click the **Add Profile** (plus) icon.
- c) In the Create/Copy Profile window that appears, enter the following to create a new Delivery Service profile:
 - **Mode:** Choose **Create** or **Copy**:
 - **Create:** To create a new profile based on the default settings for the profile, choose **Create**. You can edit these values after you create the profile.
 - **Copy:** To create a new profile based on the settings of an existing profile, choose **Copy**. You can edit these values after you create the profile.
 - **Name:** Enter a unique descriptive name for the profile. It is helpful to begin the name with **DS** to easily identify it as a Delivery Service profile.
 - **Description:** Enter a description for the new profile. This description appears when you view the list of available profiles on the Provisioning > Edit CDN > Profiles tab so enter a description that helps you identify this profile.

- **CDN:** Make sure the CDN for the Delivery Service you are configuring MSO for is selected.

Note A profile can belong to only one CDN.

- **Type:** If the mode you chose was Create, from the Type drop-down list choose **delivery_service**. If the mode you chose was Copy, you will not see this field.
- **Copy From:** If the mode you chose was Copy, from the Copy From drop-down list choose the Delivery Service profile from which to create this profile. When you create a profile by copying it from another profile, the new profile will inherit the parameters and settings of that profile. After you finish creating the new profile, you can edit these settings. If the mode you chose was Create, you will not see this field.

- d) Click **Save** to create and save the new profile.
- e) To view, and if necessary, change the settings of the new profile, choose **Provisioning > Edit CDN**.
- f) Make sure the correct CDN is selected in the CDN Name drop-down list and click the **Profiles** tab.
- g) From the Profile Details window, click **Edit** in the Action column for the new profile that you created.
- h) Click the **Multi Site Origin** tab.
- i) The Delivery Service Settings window will display the current MSO settings assigned to the profile. These settings will apply to any Delivery Service to which you assign this profile. Make changes to the fields as needed:

- **Algorithm:** Choose one of the following algorithms based on the MSO failover and load-balancing approach you want the Delivery Services that use this profile to use:
 - **strict:** With this setting, Traffic Servers will serve requests in a round robin fashion, spreading requests across multiple parents simultaneously based on the order of requests. This method is also referred to as Strict Round-Robin.
 - **true:** This setting is essentially the same as strict, but it also ensures that requests from the same IP address always go to the same Origin Server, if that server is available. This method is also referred to as IP-based Round-Robin.
 - **false:** With this setting, round robin selection does not occur. Instead, only the primary Origin Server is used as long as it is available. If the primary server is no longer available, then a backup Origin Server is used. The order in which Origin Servers are used is based on the configuration. This method is also referred to as Primary/Backup failover.
 - **consistent_hash:** Spreads requests across multiple parents simultaneously based on the hash of the content URL.
- **Parent Retry:** This setting determines whether failover of Origin Servers should be provided for 404 HTTP responses (simple retry) or 5xx HTTP responses (dead server retry):
 - **Both:** Failover is provided for both 404 HTTP responses and any 5xx HTTP responses that are configured in the “Unavailable Server Retry Response” list in the Delivery Service profile.
 - **Simple Retry:** Failover is only provided for 404 HTTP responses.
 - **Unavailable Server Retry:** Failover is only provided for the 5xx HTTP response codes that are configured in the “Unavailable Server Retry Response” list in the Delivery Service profile.
- **Unavailable Server Retry Response:** If you selected either “both” or “Unavailable Server Retry” for Parent Retry, then in this field enter the 5xx HTTP responses that should trigger failover. Only

enter one value per field. To add extra response codes, click the + icon after the last response code listed and a new field will appear. To remove one of the codes, click the – icon.

- **Max Simple Retries:** Determines the maximum number of times an Origin Server should be retried for any 404 HTTP responses.
- **Max Unavailable Server Retries:** Determines the maximum number of times an Origin Server should be retried for any of the 5xx HTTP responses configured in the Unavailable Server Retry Response setting.

Figure 3: Edit a Delivery Service Profile

- When you are finished making changes, click **Save**.
- When the profile has been successfully updated, the “Processing” status in the title bar will disappear and you will see a green check mark for every action in the Status column. Click **OK** to close the status window.

Configure the Delivery Service for MSO

Perform the following steps to configure an existing Delivery Service to use MSO:

Procedure

- Step 1** Choose **Provisioning > Edit CDN**. From the **CDN Name** drop-down list above the tabs, make sure the correct CDN is selected.

- Step 2** Click the **Delivery Services** tab.
- Step 3** From the **Delivery Service Name** drop-down list, choose the Delivery Service for which you want to configure MSO.
- Step 4** For the **Use Multi Site Origin Feature** option, choose **Yes**.
- Step 5** From the **Delivery Service Profile** drop-down list, choose the Delivery Service Profile with the MSO settings you would like to assign to this Delivery Service.
- Step 6** Assign Origin Servers to the Delivery Service. These will be used for Mid cache groups that support the Delivery Service that do not have a Parent Cache Group assigned. To assign Origin Servers, you can do one of the following:
- From the **Assign Cache Servers** drop-down list, choose the Origin Servers you would like to assign to the Delivery Service.
 - From the **Assign Device Groups to Delivery Service** drop-down list, choose the Device Groups that contain the Origin Servers that you would like to assign to the Delivery Service.
- Note** If you assign servers (Edge cache, Mid cache, or Origin) from both the Assign Cache Server list and the Assign Device Groups to the Delivery Service list, the servers that serve the Delivery Service will be a union of the two. When the MSO algorithm assigned to the Delivery Service is False, all Mid cache servers that are in this union that are not assigned a Parent Cache Group will use the rank of any Origin Servers that are in this union to determine the primary and backup Origin Servers. The Origin Server with the lowest rank is considered the primary. Also, all of the Mid caches that are in the parent cache groups of any specifically assigned Edge cache (from the “Assign Cache Servers” field) are automatically assigned to the Delivery Service.
- Note** For more information on device groups, see [Device Groups Overview, on page 181](#).
- Step 7** Click **Save**.
- Step 8** When the Delivery Service has been successfully updated, the “Processing” status in the title bar will disappear and you will see a green check mark for every action in the Status column. Click **OK** to close the status window.

What to do next

To create a new Delivery Service and configure it to support MSO, see [Add a New DNS or HTTP Delivery Service, on page 92](#).

Regex Remap Settings

By default, only the path and query string of the URL are provided for the regular expressions to match. The **Regex Remap Settings** tab of the Delivery Service profile enables you to specify parameters and settings that you can use to modify the behavior of what regular expressions will match.

To specify Regex parameters and their settings to modify the behavior of what regular expressions will match, perform the following steps:

Procedure

-
- Step 1** Choose **Provisioning > Edit CDN**.
- Step 2** Make sure the correct CDN is selected in the **CDN Name** drop-down list and click the **Delivery Services** tab.
- Step 3** To determine whether the Delivery Service already has a Delivery Service profile applied, from the Delivery Service Name drop-down list choose the Delivery Service for which to configure the Regex Remap parameters.
- Step 4** In the **Delivery Service Profile** drop-down list, determine whether a profile is currently assigned to the Delivery Service. If no profile is currently assigned to the Delivery Service, go to Step 5 to create a new Delivery Service profile. If a profile is currently assigned to the Delivery Service, perform the following steps to configure the Regex Remap settings of the profile:
- Choose **Provisioning > Edit CDN**.
 - Make sure the correct CDN is selected in the CDN Name drop-down list and click the **Profiles** tab.
 - From the Profile Details window, click **Edit** in the Action column for the profile that is assigned to the Delivery Service for which you want to change the Regex remap parameters.
 - Click the **Regex Remap Settings** tab.
 - In the **Regex Remap Parameters** field, enter the Regex remap parameters you would like to add or modify. To enter multiple parameters, separate them by a comma. You can add the following parameters:
 - **@pparam=[no-]method:** When this option is enabled, the Regex remap rules can match on the HTTP method. This option is disabled by default. To enable it, enter **@pparam=method**.
 - **@pparam=[no-]query-string:** When this option is enabled, the Regex remap rule can match against the query string. This option is enabled by default. If you would like to disable it, enter **@pparam=no-query-string**.
 - **@pparam=[no-]matrix-parameters:** When this option is enabled, the Regex remap rules can match the matrix parameters. This option is disabled by default. To enable it, enter **@pparam=matrix-parameters**.
 - **@pparam=[no-]host:** When this option is enabled, the Regex remap rules can match against the host. This option is disabled by default. To enable it, enter **@pparam=host**.
 - **@pparam=[no-]scheme:** When this option is enabled, the Regex remap rules can match against the scheme. This option is disabled by default. To enable it, enter **@pparam=scheme**.
 - When you are finished entering the Regex remap parameters, click **Save**.
 - When the profile has been successfully updated, the “Processing” status in the title bar will disappear and you will see a green check mark for every action in the Status column. Click **OK** to close the status window.

The screenshot shows the 'Provisioning > Edit CDN' interface. At the top, there's a 'CDN Name' dropdown set to 'CDN2' with a 'Delete' button. Below this is a tabbed interface with 'Servers', 'Cache Groups', 'Delivery Services', 'Client Routing', 'Profiles', and 'Content Invalidation'. The 'Profiles' tab is active, showing 'Delivery Service Settings - ds_profile'. Under this, there are three sub-tabs: 'Multi Site Origin', 'URL Signing', and 'Regex Remap Settings' (which is selected). The 'Regex Remap Parameters' field contains '@pparam=method, @pparam=host'. There are 'Cancel' and 'Save' buttons at the bottom right.

- Step 5** If no Delivery Service profile was assigned to the Delivery Service, perform the following steps to create a new Delivery Service profile to use:
- Choose **Provisioning > General Settings** and click the **Profiles** tab.
 - In the Action column header, click the **Add Profile** (plus) icon.
 - In the Create/Copy Profile window that appears, enter the following information to create a new Delivery Service profile:
 - **Mode:** Choose **Create** or **Copy**:
 - **Create:** To create a new profile based on the default settings for the profile, choose **Create**. You can edit these values after you create the profile.
 - **Copy:** To create a new profile based on the settings of an existing profile, choose **Copy**. You can edit these values after you create the profile.
 - **Name:** Enter a unique descriptive name for the profile. It is helpful to begin the name with **DS** to easily identify it as a Delivery Service profile.
 - **Description:** Enter a description for the new profile. This description appears when you view the list of available profiles on the Provisioning > Edit CDN > Profiles tab so enter a description that helps you identify this profile.
 - **CDN:** In the **CDN** drop-down list, make sure the CDN for the Delivery Service you are configuring the URL Settings for is selected.
- Note** A profile can belong to only one CDN.
- **Type:** If the mode you chose was Create, from the Type drop-down list choose **delivery_service**. If the mode you chose was Copy, you will not see this field.
 - **Copy From:** If the mode you chose was Copy, from the Copy From drop-down list choose the Delivery Service profile from which to create this profile. When you create a profile by copying it from another profile, the new profile will inherit the parameters and settings of that profile. After you finish creating the new profile, you can edit these settings. If the mode you chose was Create, you will not see this field.
 - Click **Save** to create and save the new profile.
 - Go to **Provisioning > Edit CDN**.
 - Make sure the correct CDN is selected in the CDN Name drop-down list and click the **Profiles** tab.
 - From the Profile Details window, click **Edit** in the Action column for the new profile that you created.

- h) In the Regex Remap Parameters field, verify or enter the Regex remap parameters you would like to add or modify. To enter multiple parameters, separate them by a comma. You can add the following parameters:
- **@pparam=[no-]method:** When this option is enabled, the Regex remap rules can match on the HTTP method. This option is disabled by default. To enable it, enter **@pparam=method**.
 - **@pparam=[no-]query-string:** When this option is enabled, the Regex remap rule can match against the query string. This option is enabled by default. If you would like to disable it, enter **@pparam=no-query-string**.
 - **@pparam=[no-]matrix-parameters:** When this option is enabled, the Regex remap rules can match the matrix parameters. This option is disabled by default. To enable it, enter **@pparam=matrix-parameters**.
 - **@pparam=[no-]host:** When this option is enabled, the Regex remap rules can match against the host. This option is disabled by default. To enable it, enter **@pparam=host**.
 - **@pparam=[no-]scheme:** When this option is enabled, the Regex remap rules can match against the scheme. This option is disabled by default. To enable it, enter **@pparam=scheme**.
- i) If you made any changes to the Regex remap parameters, click **Save**.
- j) When the profile has been successfully updated, the “Processing” status in the title bar will disappear and you will see a green check mark for every action in the Status column. Click **OK** to close the status window.

Step 6

If you created a new Delivery Service profile to use for the Regex Remap settings, perform the following steps to assign the new profile to the Delivery Service:

- a) Choose **Provisioning > Edit CDN**.
- b) Make sure the correct CDN is selected in the CDN Name drop-down list and click the **Delivery Services** tab.
- c) From the **Delivery Service Name** drop-down list, choose the Delivery Service for which you are configuring the Regex Remap settings.
- d) From the **Delivery Service Profile** drop-down list, choose the new profile that you created.
- e) Click **Save** to save the changes.

Edge Geo Blocking

Without Edge Geo Blocking configured on the Delivery Service, for Delivery Services that use DNS routing the Traffic Router never sees the complete URL requested by the client or the source IP address of the client making the request. Because the Traffic Router never sees the source IP address of the client making the request, Delivery Services that use DNS routing have the following limitations:

- For Delivery Services that use CZF or NGB *and* DNS routing, the Traffic Router uses the IP address of the DNS resolver and *not* the IP address of the actual client to determine whether to allow the request.
- Anonymous blocking and ASN blocking are not supported.



Note These limitations do *not* apply to Delivery Services that use HTTP-based routing because for HTTP-based routing requests, the Traffic Router will see the actual client IP address and full URL of the request.

The Edge Geo Blocking feature enables the Edge caches to participate in determining whether a client request should be allowed. With this feature enabled, the Edge cache passes the IP address of the actual client making the request and the requested URL to the Traffic Router to determine whether the request should be allowed. Therefore, Delivery Services that use DNS routing that have Edge Geo Blocking enabled provide the following:

- Support for anonymous blocking
- Support for ASN blocking
- For all of the blocking mechanisms (CZF, NGB, Anonymous blocking, and ASN blocking), the Traffic Router can use the IP address of the actual client to determine whether to allow the request.

Because the Edge caches now participate in determining whether a client request should be allowed, Edge Geo Blocking also ensures that clients cannot bypass blocking policies for HTTP-based or DNS-based Delivery Services by bypassing the Traffic Router.

When Edge Geo Blocking is enabled, the Traffic Router maintains all existing blocking functionality and still performs the initial check on the client request. The Edge caches have no autonomy to make blocking decisions. Edge caches can only implement the blocking instructions provided by the Traffic Router.

The Edge cache will check with the Traffic Router only for requests handled by a Delivery Service that has Edge Geo Blocking enabled *and* has a client blocking option configured (CZF, NGB, Anonymous blocking, or ASN blocking).



Note Any changes to the blocking policy will impact in-progress sessions.

Configure Edge Geo Blocking

The Delivery Service profile is used to enable and configure the Edge Geo Blocking feature for a Delivery Service. By default, the Edge Geo Blocking feature is disabled.

Perform the following steps to configure a Delivery Service profile for URL Signing:



Note Any changes to the blocking policy will impact in-progress sessions.

Procedure

- Step 1** Choose **Provisioning > Edit CDN**.
- Step 2** Make sure the correct CDN is selected in the CDN Name drop-down list and click the **Delivery Services** tab.
- Step 3** To determine whether the Delivery Service already has a Delivery Service profile applied, from the Delivery Service Name drop-down list choose the Delivery Service to configure for Edge Geo Blocking.

- Step 4** In the Delivery Service Profile drop-down list, determine whether a profile is currently assigned to the Delivery Service. If no profile is currently assigned to the Delivery Service, go to Step 5 to create a new Delivery Service profile. If a profile is currently assigned to the Delivery Service, perform the following steps to configure the Edge Geo Blocking settings of the profile:
- Choose **Provisioning** > **Edit CDN**.
 - Make sure the correct CDN is selected in the CDN Name drop-down list and click the **Profiles** tab.
 - From the Profile Details window, click **Edit** in the Action column for the profile that is assigned to the Delivery Service for which you are configuring Edge Geo Blocking.
 - Click the **Edge Blocking** tab.
 - The Delivery Service Settings window will display the current Edge Geo Blocking settings assigned to the profile. These settings will apply to any Delivery Service to which you assign this profile. Configure the following fields:
 - **Enable Edge Geo Blocking:** Check this check box to enable the Edge Geo Blocking feature.
 - **TTL (seconds):** Enter the number of seconds the Allow/Deny response from the Traffic Router should remain in the authproxy cache. The default is 60 seconds.
 - When you are finished making changes, click **Save**.
 - When the profile has been successfully updated, the “Processing” status in the title bar will disappear and you will see a green check mark for every action in the Status column. Click **OK** to close the status window.

The screenshot shows the 'Provisioning > Edit CDN' interface. At the top, there's a 'CDN Name' dropdown set to 'cdn' and a 'Delete' button. Below this is a tabbed interface with tabs for 'Servers', 'Cache Groups', 'Delivery Services', 'Client Routing', 'Profiles', and 'Content Invalidation'. The 'Profiles' tab is active, showing 'Delivery Service Settings - DS_Geo1'. Under this, there are sub-tabs: 'Multi Site Origin', 'URL Signing', 'Regex Remap Settings', 'CDN Routing', and 'Edge Blocking'. The 'Edge Blocking' sub-tab is selected, displaying a checked checkbox for 'Enable Edge Geo Blocking' and a text input field for 'TTL (seconds)' with the value '60'. At the bottom right of the settings area are 'Cancel' and 'Save' buttons.

- Step 5** If no Delivery Service profile was assigned to the Delivery Service, perform the following steps to create a new Delivery Service profile to use:
- Choose **Provisioning** > **General Settings** and click the **Profiles** tab.
 - In the Action column header, click the **Add Profile** (plus) icon.
 - In the Create/Copy Profile window that appears, enter the following information to create a new Delivery Service profile:

- **Mode:** Choose **Create** or **Copy**:
 - **Create:** To create a new profile based on the default settings for the profile, choose **Create**. You can edit these values after you create the profile.
 - **Copy:** To create a new profile based on the settings of an existing profile, choose **Copy**. You can edit these values after you create the profile.
 - **Name:** Enter a unique descriptive name for the profile. It is helpful to begin the name with **DS** to easily identify it as a Delivery Service profile.
 - **Description:** Enter a description for the new profile. This description appears when you view the list of available profiles on the Provisioning > Edit CDN > Profiles tab so enter a description that helps you identify this profile.
 - **CDN:** In the CDN drop-down list, make sure the CDN for the Delivery Service you are configuring the URL Settings for is selected.
- Note** A profile can belong to only one CDN.
- **Type:** If the mode you chose was Create, from the Type drop-down list choose **delivery_service**. If the mode you chose was Copy, you will not see this field.
 - **Copy From:** If the mode you chose was Copy, from the Copy From drop-down list choose the Delivery Service profile from which to create this profile. When you create a profile by copying it from another profile, the new profile will inherit the parameters and settings of that profile. After you finish creating the new profile, you can edit these settings. If the mode you chose was Create, you will not see this field.
- d) Click **Save** to create and save the new profile.
 - e) Go to **Provisioning > Edit CDN**.
 - f) Make sure the correct CDN is selected in the CDN Name drop-down list and click the **Profiles** tab.
 - g) From the Profile Details window, click **Edit** in the Action column for the new profile that you created.
 - h) Click the **Edge Blocking** tab.
 - i) The Delivery Service Settings window will display the current Edge Geo Blocking settings assigned to the profile. These settings will apply to any Delivery Service to which you assign this profile. Verify or configure the following fields:
 - **Enable Edge Geo Blocking:** Check this check box to enable the Edge Geo Blocking feature.
 - **TTL (seconds):** Enter the number of seconds the Allow/Deny response from the Traffic Router should remain in the authproxy cache. The default is 60 seconds.
 - j) If you made any changes, click **Save**.
 - k) When the profile has been successfully updated, the “Processing” status in the title bar will disappear and you will see a green check mark for every action in the Status column. Click **OK** to close the status window.
- Step 6** If you created a new Delivery Service profile to use for Edge Geo Blocking, perform the following steps to assign the new profile to the Delivery Service:
- a) Choose **Provisioning > Edit CDN**.
 - b) Make sure the correct CDN is selected in the **CDN Name** drop-down list and click the **Delivery Services** tab.

- c) From the **Delivery Service Name** drop-down list, choose the Delivery Service for which you are configuring Edge Geo Blocking.
- d) From the **Delivery Service Profile** drop-down list, choose the new profile that you created.
- e) Click **Save** to save the changes.

CDN Routing Name

By default when the Traffic Router sends an HTTP 302 redirect message to the requesting client, the FQDN in the response has the following format:

`<cache_hostname>.<ds_name>.<cdn_domain_name>`

where: `<cache_hostname>` is the hostname of the edge cache server to which the client is being redirected, `<ds_name>` is the name of the Delivery Service that handles the content being requested, and `<cdn_domain_name>` is the full domain name of the CDN. For example, `http://edge1.vod-west.cdn.companyx.com/movie/master.m3u8`

The CDN Routing Name setting enables you to insert a subdomain into the default FQDN. This new subdomain will be treated as a new zone for DNS. To specify a string to insert into the default FQDN used for the HTTP 302 redirect messages, perform the following steps:

Procedure

- Step 1** Choose **Provisioning > Edit CDN** and click the **Profiles** tab.
- Step 2** Click **Edit** in the action column for the Delivery Service profile for which you want to insert a subdomain. All HTTP 302 redirect messages that the Traffic Router sends for any requests that are serviced by a Delivery Service that is assigned this profile will insert the string you specify.
- Step 3** Click the **CDN Routing** tab.
- Step 4** In the **CDN Routing Name** field, enter the subdomain that you would like to insert into the HTTP 302 redirect messages that the Traffic Router sends for any requests that are serviced by a Delivery Service that is assigned this profile.

The screenshot shows the 'Provisioning > Edit CDN' interface. At the top, there's a 'CDN Name' dropdown set to 'omd-cdn1' with a 'Delete' button. Below this are tabs for 'Servers', 'Cache Groups', 'Delivery Services', 'Client Routing', 'Profiles', and 'Content Invalidation'. The 'Profiles' tab is active, showing 'Delivery Service Settings - ds-profile'. Under this, there are sub-tabs: 'Multi Site Origin', 'URL Signing', 'Regex Remap Settings', and 'CDN Routing'. The 'CDN Routing' sub-tab is selected, displaying a 'CDN Routing Name' field with the value 'OMD'. At the bottom right of this section are 'Cancel' and 'Save' buttons.

- Step 5** Click **Save**.



CHAPTER 7

Manage Client Routing

A Traffic Router handles client requests for content and determines which client requests to allow. After it has determined that a client request is allowed, the Traffic Router uses client routing methods to determine which cache group is the best group to deliver the content for a client request. This chapter describes the client routing methods that a Traffic Router can use to determine which cache group is the best to use for a request.



Note For more information on how a Traffic Router determines which client requests are allowed, see the [Understanding CDN Client Blocking Options](#) section in the "Initial CDN Provisioning" chapter.

This chapter includes the following topics:

- [Client Routing Overview, on page 155](#)
- [Coverage Zone File, on page 156](#)
- [Proximity Routing, on page 159](#)
- [Geolocation Based Routing, on page 161](#)

Client Routing Overview



Note If you are running OMD Director in an HA configuration, you can only make configuration changes from the Primary OMD Director instance. To perform the steps in this section, make sure you are logging into the Primary OMD Director instance.

The Traffic Router supports the following client routing methods to determine which cache group is the best to use for a client request:

- **Coverage Zone File (CZF):** The CZF is a static JSON file that maps IP address ranges to cache groups. The Traffic Router checks the CZF for an IP address range that matches the requesting address to determine the best cache group to deliver the content.
- **Proximity Routing:** Proximity routing uses network proximity maps that leverage routing information databases to help determine the best cache group to deliver the content.



Note Currently Proximity routing requires the use of VDS-IS Proximity Engines (PxE) and is intended for OMD customers that are migrating from VDS-IS.



Note Currently the Proximity Server does not support the rating of IPv6 PSAs or PTAs.

- **Geolocation Based Routing:** Geolocation based routing uses a geolocation database to determine the best cache group to deliver the content.

The Traffic Router uses these client routing methods in the following order to determine the best cache to redirect a client request to:

1. The Traffic Router checks the CZF file first. (See [Coverage Zone File](#), on page 156 for more details.)
2. If the check of the CZF file does not return a Cache to use for the request, Proximity- routing is checked next, if it is enabled in the profile of the Traffic Router. (See [Proximity Routing](#), on page 159 for more details.)
3. If Proximity routing is not enabled or if Proximity routing does not return a Cache to use for the request, the Traffic Router will use geolocation based routing to determine a cache group the client request should use and then chooses a cache from that group. If no geolocation is available for the requesting address, the Geo Miss Default Latitude and Geo Miss Default Longitude settings of the Delivery Service are used to determine the closest cache group to use for the client request and then chooses a cache from that cache group. (See [Geolocation Based Routing](#), on page 161 for more details.)
4. If no Caches are returned based on geolocation routing, the client will receive a “503 Service Unavailable” message.

Coverage Zone File

The CZF file is a static JSON file that maps IP address ranges to cache groups. The Traffic Router uses the CZF file to help determine which cache to redirect a client request to. To determine which cache to redirect a client request to, the Traffic Router performs the following steps:

1. The Traffic Router checks the CZF file for a cache group with a subnet match to one of the following addresses:
 - The client IP address making the request, for HTTP-based Delivery Services
 - The DNS Resolver IP address making the request, for DNS-based Delivery Services that do not have the Enhanced DNS Request Routing (ECS) option enabled in the profile of the Traffic Router or if the DNS query has no ECS option
 - The subnet in the EDNS0 option field, which is the subnet of the client IP address making the request, for DNS-based Delivery Services that have the Enhanced DNS Request Routing (ECS) option enabled on the profile of the Traffic Router and the DNS query has an ECS Option



Note For information on the Enhanced DNS Request Routing (ECS) feature and how to configure it, see [Configuring Enhanced DNS Request Routing, on page 271](#).



Note When DNS-based Delivery Services are configured, the Coverage Zone file needs to have entries with respect to the IP address of the DNS proxies. If Edge Geo Blocking is enabled on the Delivery Service, the Coverage Zone file must also have entries for the actual client IP address.

2. If a cache group is found, the Traffic Router then chooses a Cache from that group based on cache availability, cache load, and cache content.
3. If there is no available cache for the cache group that was matched, the Traffic Router performs the following steps to look for a backup Edge cache group:
 1. The Traffic Router checks to see if the *original* Edge cache group has any fallback cache groups assigned. If the original Edge cache group has fallback cache groups assigned, the Traffic Router will check this list of fallback cache groups *in the order listed*, to find an available cache.
 2. If no fallback cache groups are configured on the original Edge cache group that was matched, or if there are no Edge caches available from the fallback cache groups, the Traffic Router checks the "Fallback Enable" setting of the original Edge cache group that was matched:
 1. If the Fallback Enable setting is set to **No**, no further searches are performed and no cache is returned. The client is redirected to the bypass location listed in the Bypass FQDN field of the Delivery Service. If the Bypass FQDN field is empty, the client will receive a 503 Service Unavailable message.
 2. If the Fallback Enable setting is set to **Yes**, the Traffic Router looks for the next closest cache group to the original group, based on the latitude and longitude configured in the CZF file. The Traffic Router then chooses an Edge cache from that group based on cache availability, cache load, and cache content.
 3. When the Geo Limit setting is "CZF Only", if there is no cache available from the next closest group, or if there are no latitude and longitude settings in the CZF file for any of the other caches assigned to the Delivery Service for the client request, the client is redirected to the bypass location listed in the Bypass FQDN field of the Delivery Service. If the Bypass FQDN field is empty, the client will receive a 503 Service Unavailable message. If the Geo Limit setting is *not* set to "CZF Only", the Traffic Router goes on to the next step.



Note For more information on backup Edge cache groups, including how to configure them, see [Backup Edge Cache Groups, on page 78](#).

4. If the CZF check does not return a cache, the Traffic Router will next use Proximity routing to determine which Cache to use for the request. (See [Proximity Routing, on page 159](#) for more details.) If Proximity

routing is not enabled, the Traffic Router will use geolocation based routing to determine which Cache to use for the request. (See [Geolocation Based Routing, on page 161](#) for more details.)

**Note**

When Geo Limit is set to "CZF Only" or "CZF + CountryCodes(s)", the CZF file is also used to help determine which client requests to allow. For more information on how the CZF is used with client blocking, see [Understanding CDN Client Blocking Options, on page 23](#).

To assign a CZF file to a Delivery Service, refer to [Managing the CZF File, on page 158](#). To see an example of a CZF file and its format, see [Example CZF File, on page 257](#).

Managing the CZF File

You can see which CZF file the Traffic Routers are using for a specific CDN and view its contents from the following locations:

- **Provisioning > Edit CDN > Client Routing**

Provisioning > Edit CDN

CDN Name:

Coverage Zone File

Name	Description	Active	Creation Date	Action
Tl.json	used for TI	false	2017-11-24T07:18:19.346464Z	View

Showing 1 to 1 of 1 entries

- **Provisioning > General Settings > Coverage Zone**

Provisioning > General Settings

General Settings

Anonymous IP Database Profiles **Coverage Zone** Device Groups

Coverage Zone File

Name	Description	CDN	Creation Date	Action
cz1.json				View Delete
Tl.json	used for TI	CDN1	2017-11-24T07:18:19.346464Z	View Delete

Showing 1 to 2 of 2 entries

Upload a CZF File

From Coverage Zone window you can also upload a new CZF file. To upload a new CZF file to the system, perform the following steps:

Procedure

Step 1 From the **Provisioning > General Settings > Coverage Zone** window, in the **Action** column, click the **Add CZF** icon (plus sign).

Step 2 In the Add CZF window that appears, click **Browse** and browse for the CZF file to upload. When you find the file, select it and click **Open**.

Note To see an example of a CZF file and its format, see [Example CZF File, on page 257](#).

Step 3 Optionally enter a description for the CZF file.

Step 4 From the CDN drop-down list, choose the CDN that will use this CZF file. All Traffic Routers that belong to the CDN that you choose will use this CZF file.

Step 5 Click **Save** to save the file.

Step 6 After you have added the CZF file, you need to assign the CZF file to the Traffic Router profiles that should use it. To assign a CZF file to a Traffic Router profile, see [Assign a CZF to a Traffic Router Profile, on page 166](#).

Delete a CZF File

From the Provisioning > General Settings > Coverage Zone window you can also delete a CZF file. To delete a CZF file, in the row for the CZF file that you want to delete, click **Delete** in the Action column.

Proximity Routing

Proximity routing enables the Traffic Router to more intelligently choose which cache group to redirect the client request to by using the network proximity of the cache group. Traffic Routers use external Proximity Servers to help determine the network proximity of the cache group. To calculate the cache group's network proximity, the Proximity Servers run routing protocols to receive routing updates from network routers. OMD uses existing VDS-IS Proximity Engines (PxE) as the network Proximity Servers.



Note Currently Proximity routing requires the use of VDS-IS Proximity Engines (PxE) and is intended for OMD customers that are migrating from VDS-IS.

A Proximity Server can listen for OSPF, BGP, and IS-IS updates (depending on its configuration) and provide proximity information between clients requesting content and the cache groups that can deliver the content. The Proximity Server leverages routing information databases (IGP and BGP) by interconnecting and peering with routers in the network. The Proximity Server uses this information to provide a list of Caches to the Traffic Router, ranked in order of optimal routes for content and messages in the network based on the proximity rating of the Caches.

If more than one Proximity Server is configured, then the Traffic Router will send a rating request to the first available Proximity Server to “rate” the Proximity Servers. From the response, the Traffic Router will have the ratings for all configured Proximity Servers. Using this information, the Traffic Router will be able to send the proximity requests for the client requests to the closest Proximity Server. This is referred as the PxE rating procedure.

When the Traffic Router uses Proximity routing, the following steps occur to determine which Cache to use for a client request:

1. The Traffic Router sends a cache rating request, using a Proximity API, to the closest Proximity Server based on the Proximity Server rating. This request includes the following:
 - **Proximity Source Address (PSA):** This is the IP address that is used as the root of the calculation. The proximity function will have to find the closest address to the PSA. Which address is used for the PSA is based on the following:
 - The client IP address making the request, for HTTP-based Delivery Services
 - The DNS Resolver IP address making the request, for DNS-based Delivery Services that do not have the Enhanced DNS Request Routing (ECS) option enabled on the profile of the Traffic Router or if the DNS query has no ECS option
 - The subnet in the EDNS0 option field, which is the subnet of the client IP address making the request, for DNS-based Delivery Services that have the Enhanced DNS Request Routing (ECS) option enabled on the profile of the Traffic Router and the DNS query has an ECS Option.



Note For information on the Enhanced DNS Request Routing (ECS) feature and how to configure it, see [Configuring Enhanced DNS Request Routing, on page 271](#).

- **Proximity Target List (PTL):** This list contains two or more PTA IP addresses, which are the IP addresses of Edge Caches. The Proximity Algorithm will use these IP addresses to determine which IP address is closest to the PSA.
 - **Proximity Ranking Depth:** This is an integer number that determines the length of the ranking list that the Proximity Server should return.
2. Using the information in the caching request from the Traffic Router, the Proximity Server calculates the network distance between the PSA and the PTAs in the TPL, to calculate the proximity rating of each PTA.

3. The Proximity Server returns a list of the PTAs with their proximity rating to the Traffic Router.
4. The Traffic Router chooses the cache group with the best proximity rating, and then chooses a Cache to use from that group based on cache availability, cache load, and cache content.
5. If Proximity routing does not return a cache to use or there is any error during the Proximity routing process, the Traffic Router will use geolocation based routing to choose a cache to handle the client request. (See [Geolocation Based Routing, on page 161](#) for more details.)

See [Configure Proximity Routing, on page 166](#) for steps to configure Proximity routing.

Geolocation Based Routing

**Note**

Prior to using Geolocation-based routing, you should ensure that the MaxMind geolocation database has been configured in your Media Streamer deployment. For details on performing this operation, refer to the "Configure the MaxMind Database for Geolocation" appendix in the *Cisco Media Streamer and Cisco Media Broadcaster Installation and Upgrade Guide*.

When the Traffic Router uses Geolocation-based routing, the following steps occur to determine which cache to use for a client request:

1. The Traffic Router uses a geolocation database to determine the latitude and longitude of the requesting address, where the requesting address is one of the following:
 - The client IP address making the request, for HTTP-based Delivery Services
 - The DNS Resolver IP address making the request, for DNS-based Delivery Services that do not have the Enhanced DNS Request Routing (ECS) option enabled on the profile of the Traffic Router or if the DNS query has no ECS option
 - The subnet in the EDNS0 option field, which is the subnet of the client IP address making the request, for DNS-based Delivery Services that have the Enhanced DNS Request Routing (ECS) option enabled on the profile of the Traffic Router and the DNS query has an ECS Option.

**Note**

For information on the Enhanced DNS Request Routing (ECS) feature and how to configure it, see [Configuring Enhanced DNS Request Routing, on page 271](#).

2. The Traffic Router compares the latitude and longitude configured for each cache group with the latitude and longitude determined in Step 1 to find the geographically closest cache group. The Traffic Router then chooses a Cache to use from that group based on cache availability, cache load, and cache content.
3. If no geolocation is available, the Geo Miss Default Latitude and Geo Miss Default Longitude settings of the Delivery Service are used to determine the closest cache group to use. The Traffic Router then chooses a Cache to use from that group based on cache availability, cache load, and cache content.



CHAPTER 8

Manage Profiles

The Profiles tab of the Edit CDN page lists the profiles that are available in the CDN, displays the type of each profile, and displays which servers in the CDN are using the profile. This chapter describes the profiles available from this tab, their settings, and how to configure them.

This chapter includes the following topics:

- [Profiles Overview, on page 163](#)
- [Add a Profile, on page 165](#)
- [Assign a CZF to a Traffic Router Profile, on page 166](#)
- [Configure Proximity Routing, on page 166](#)
- [Anonymous Blocking, on page 169](#)
- [ASN Blocking, on page 173](#)
- [NGB Whitelist, on page 176](#)

Profiles Overview

Some of the features that are supported by Delivery Services, Traffic Routers, and Origin Servers must be configured using profiles. From the Profiles tab of the Provisioning CDN > Edit CDN window you can edit profiles that are needed for this purpose. From the Profiles tab, you can configure the following types of profiles and settings:

- **Traffic Router profiles:** Traffic Router profiles contain the following settings:
 - **Proximity Routing settings:** Proximity routing enables the Traffic Router to more intelligently choose which cache group to redirect the client request to by using the network proximity of the cache group. For information on how the Traffic Router uses Proximity routing and how to configure it, see [Proximity Routing, on page 159](#).
 - **Anonymous Blocking settings:** Anonymous Blocking uses the MaxMind Anonymous IP database to identify requests that are coming from commercial VPN services, Tor Exit Nodes, Hosting Providers, and Public Proxies. You can configure the Media Streamer deployment to use this information to block traffic from these types of sources. For information on how the Traffic Router uses Anonymous Blocking and how to configure it, see [Anonymous Blocking, on page 169](#).
 - **Coverage Zone File:** The Traffic Router uses the CZF file to help determine which cache to redirect a client request to. When GeoLimit is set to "CZF Only" or "CZF + CountryCodes(s)", the Traffic Router also uses the CZF file to help determine which client requests to allow. For more information on how the Traffic Router uses the CZF file, see [Coverage Zone File, on page 156](#). For information

on how to specify the coverage zone file this Traffic Router profile will use, see [Assign a CZF to a Traffic Router Profile, on page 166](#).

- **NGB Whitelist:** An NGB white list is an optional whitelist that can be configured to work with the National geoblocking feature to identify addresses that National geoblocking should not block. The NGB whitelist is only used when Geo Limit is set to "CZF + CountryCode(s)". For information on how to configure an NGB whitelist, see [NGB Whitelist, on page 176](#).
- **ISP Database:** ASN blocking uses the MaxMind ISP or the GeoLite "ASN" database to determine the AS number of an IP address. From the ISP Database tab in the Traffic Router profile, you can identify which database file ASN blocking should use. For information on how to upload this database and to configure ASN blocking, see [ASN Blocking, on page 173](#).

- **Delivery Service profiles:** Delivery Service profiles contain the following settings:

- **Multi Site Origin settings:** The MSO feature enables the Mid cache servers to be aware that there are multiple Origin Servers available to serve the Delivery Service, which enables them to provide Origin Server failover and load-distribution features. For information on the Multi Site Origin feature and how to configure it, see [Multi Site Origin, on page 133](#).
- **URL Signing:** The URL signing feature of Media Streamer provides the infrastructure to validate content URLs to prevent unauthorized access. For information on the URL Signing feature and how to configure it, see [URL Signing, on page 127](#).
- **Regex Remap settings:** By default, only the path and query string of the URL are provided for the regular expressions remap rules to match. The Regex Remap Settings tab of the Delivery Service profile enables you to specify parameters and settings that you can use to modify the behavior of what regular expressions will match. For more information on how to configure the Regex Remap parameters, see [Regex Remap Settings, on page 147](#).
- **CDN Routing:** By default when the Traffic Router sends an HTTP 302 redirect message to the requesting client, the FQDN in the response has the following format:

`<cache_hostname>.<ds_name>.<cdn_domain_name>`

where: `<cache_hostname>` is the hostname of the edge cache server to which the client is being redirected, `<ds_name>` is the name of the Delivery Service that handles the content being requested, and `<cdn_domain_name>` is the full domain name of the CDN. For example, `http://edge1.vod-west.cdn.companyx.com/movie/master.m3u8`

The **CDN Routing** tab of the Delivery Service profile enables you to identify a string to insert into the default FQDN to identify this as a CDN response. For more information on how to configure this string, see **CDN Routing Name** section in [Manage Delivery Services, on page 91](#).

- **Edge Blocking:** The Edge Geo Blocking feature enables Delivery Services that use DNS routing to support Anonymous blocking and ASN blocking. It also allows Delivery Services that use DNS routing and CZF or NGB to use the IP address of the actual client instead of the IP address of the DNS resolver to determine whether to allow the request. For information on the Edge Geo Blocking feature, including how to configure it, see [Edge Geo Blocking, on page 150](#).
- **Origin Server profiles:** Origin Server profiles contain the following settings:
 - **Multi Site Origin settings:** The MSO feature enables the Mid cache servers to be aware that there are multiple Origin Servers available to serve the Delivery Service, which enables them to provide Origin Server failover and load-distribution features. For information on the Multi Site Origin feature and how to configure it, see [Multi Site Origin, on page 133](#).



Note If you are running OMD Director in an HA configuration, you can only make configuration changes from the Primary OMD Director instance. To perform the steps in this section, make sure you are logging into the Primary OMD Director instance.

Add a Profile

If a profile does not exist that contains the settings that you need, perform the following steps to create a new profile:

Procedure

- Step 1** Choose **Provisioning > General Settings** and click the **Profiles** tab.
- Step 2** In the Action column header, click the **Add Profile** (plus) icon.
- Step 3** In the Create/Copy Profile window that appears, configure the following fields to create a new profile:
- **Mode:** Choose **Create** or **Copy**:
 - **Create:** To create a new profile based on the default settings for the profile, choose **Create**.
 - **Copy:** To create a new profile based on the settings of an existing profile, choose **Copy**.
- Note** For Origin and Delivery Service profiles you can choose either **Create** or **Copy**. For all other profile types you can only choose **Copy** to create a new profile.
- **Name:** Enter a unique descriptive name for the profile. It is helpful to begin the name with letters that help identify what type of profile it is, such as DS_ for a Delivery Service profile, or ORG_ for an Origin Server profile.
 - **Description:** Enter a description for the new profile. This description appears when you view the list of available profiles on the Provisioning > Edit CDN > Profiles tab so enter a description that helps you identify this profile.
 - **CDN:** From the CDN drop-down list, choose the CDN to which this profile should belong.
- Note** A profile can belong to only one CDN.
- **Type:** If you chose to create a new profile, choose either **Origin** or **Delivery_Service** depending on the type of profile you are creating.
- Note** You will not see the Type field if you chose Copy as the mode to add the new profile.
- **Copy From:** If you chose to add the new profile by copying an existing profile, choose the profile from which to create this profile. When you create a profile by copying from another profile, the new profile will inherit the settings of that profile. After you finish creating the new profile, you can edit these settings.
- Note** You will not see the Copy From field if you chose Create as the mode to add the new profile.

Step 4 Click **Save** to create and save the new profile.

Assign a CZF to a Traffic Router Profile

To allow Traffic Routers in the same CDN to use different CZF files, you can assign different CZF files to different Traffic Router profiles. Traffic Routers will use the CZF file that is assigned to the Traffic Router profile to which they are assigned.

Perform the following steps to assign a CZF file to a Traffic Router profile:

Procedure

- Step 1** Choose **Provisioning** > **Edit CDN** and click the **Profiles** tab.
- Step 2** Click **Edit** in the action column for the Traffic Router profile for which you want to assign the CZF file.
- Step 3** Click the **Coverage Zone File** tab.
- Step 4** From the **Coverage Zone File** drop-down list, choose the CZF file to assign and click **Save**.

Note You can only assign CZF files that have been uploaded to the CDN. For information on how to upload a CZF file to the CDN, see [Managing the CZF File, on page 158](#).

Configure Proximity Routing

To use Proximity routing, you need to configure every profile that is used by a Traffic Router that should use Proximity routing. Perform the following steps to configure a Traffic Router profile for Proximity routing:

Procedure

- Step 1** From OMD Director, choose **Provisioning** > **Edit CDN**. From the CDN Name drop-down list, choose the CDN that contains the Traffic Router profile for which you want to configure Proximity routing.
- Step 2** Click the **Profiles** tab and from the Profiles Detail section, click **Edit** for the Traffic Router profile for which you are configuring Proximity routing.

Note All of the Traffic Routers in a CDN must be assigned the same Traffic Router profile.

- Step 3** From the Routing Settings pane that appears, click the **Proximity Routing** tab.
- Step 4** On the page that appears, use the Proximity Routing Parameters table to configure the parameters. When finished entering the parameters, click **Save** to save the changes.

Table 1: Proximity Routing Parameters

Parameter Name	Settings	Default Value	Description
Enable Proximity Routing	Enabled (checked) or Disabled (unchecked)	Disabled	This parameter determines whether Network Proximity based routing is enabled or disabled.
Distributed Among Closest	Enabled (checked) or Disabled (unchecked)	Disabled	When Disabled, the Traffic Router will always try the first Proximity Server with the lowest PxE rating, based on the PxE rating procedure, discussed in the Proximity Routing section. If there is a load balancer deployed in front of the Proximity Servers, this parameter should be Disabled.
Dead Retry Interval	0-86400 (seconds)	60 seconds	This parameter determines the number of seconds before a failed request to a Proximity Server is retried. Valid values: 0-86000 seconds.
Ranking Depth	1-1024	60	Determines the length of the PTA ranking list to be returned by a Proximity Server.
Cache Timeout	0-86400 (seconds)	1800 (seconds)	The Traffic Router caches the proximity results and will reuse the results for all other subsequent requests from the same subnet to improve performance. This parameter determines how long the proximity results from the Proximity Servers can be cached. A value of 0 indicates that proximity results are not cached.
Maximum Cache Entries	1-16777216	1048576	This parameter determines the maximum number of entries that can be cached.
Proximity Servers: Host	Enter the IP address of a Proximity Server to use	N/A	The list of hosts determines the Proximity Servers to use. Note Currently the Proximity Servers that OMD uses must be VDS-IS Proximity Engines (PxE). Network Proximity is intended for OMD customers that are migrating from VDS-IS. Note Network Proximity based routing is not supported for IPv6 addresses. So the IP address of the Proximity Server must be a valid IPv4 address.

Parameter Name	Settings	Default Value	Description
Proximity Servers: Password	Currently passwords are not supported by OMD Release 3.10. If the existing Proximity Servers are configured with a password, please remove the password using the VDS-IS CDSM or Service Router CLI.		

Example

The following shows an example of configuring Proximity routing:

Routing Settings - CCR_CDN

Proximity Routing
Anonymous Blocking

Status success

Enable Proximity Routing ☒

Distributed Among Closest ☒

Dead Retry Interval * 60

Ranking Depth * 60

Cache Timeout * 1800

Maximum Cache Entries * 86400

Proximity Servers:

Host	Password	
10.63.231.100		-
10.63.231.101		- +

Cancel Save

Configure the Proximity Server

For information on how to configure the settings of the Proximity Server, refer to the Cisco Videoscape Distribution Suite, Internet Streamer Software Configuration Guide. For Cisco VDS-IS 4.3.2, you can find this information at:

https://www.cisco.com/c/en/us/td/docs/video/cds/cda/is/4_3_2/configuration_guide/SCG1/configdevice.html#35860.

For other versions of Cisco VDS-IS, refer to the following page for links to the different configuration guides:

<https://www.cisco.com/c/en/us/support/video/videoscape-distribution-suite-internet-streaming/products-installation-and-configuration-guides-list.html>.

Anonymous Blocking



Note Anonymous Blocking is not supported by DNS Delivery Services unless Edge Geo Blocking is configured. For more information on Edge Geo Blocking, see [Edge Geo Blocking, on page 26](#).

Anonymous Blocking uses the MaxMind Anonymous IP database to identify requests that are coming from commercial VPN services, Tor Exit Nodes, Hosting Providers, and Public Proxies. You can configure the OMD deployment to use this information to block traffic from these types of sources. You do not have to block traffic from all four types of sources; you can choose which of the four types of sources to block. When you enable the Anonymous Blocking feature, you can also create white lists for IPv4 and IPv6 addresses to identify IP addresses that you want ensure Anonymous Blocking does *not* block.

When the profile for the Traffic Router is configured for Anonymous Blocking, when traffic is received for a Delivery Service that is not blocked by the CZF file or National geoblocking, the Traffic Router will check to see if the Delivery Service has Anonymous Blocking enabled. If the Delivery Service has Anonymous Blocking enabled, the Traffic Router will first check the IP address of the request against the white lists. If the request matches an entry in the white list, the request is allowed.

If the IP address of the request does not match an address in the white list, the Traffic Router will check if the IP address is in the Anonymous IP Database. If the IP address is not in the database, the client is allowed. If the IP address is in the database, the database will determine what anonymizer type the request is from, Anonymous VPN, Hosting Provider, Public Proxy, or Tor Exit Node. If the request is from an anonymizer type that you have configured the Traffic Router profile to block, the Traffic Router will block the request. Otherwise, the Traffic Router will allow the request.

Requests that are blocked by Anonymous Blocking will be redirected to the Anonymous Blocking Redirect URL, if configured. Otherwise the client will receive a 403 Forbidden message.

Configure Anonymous Blocking



Note Anonymous Blocking is not supported by DNS Delivery Services unless Edge Geo Blocking is configured. For more information on Edge Geo Blocking, see [Edge Geo Blocking, on page 26](#).

To configure Anonymous Blocking, you must do the following:

1. Upload an anonymous IP database to use for blocking.
2. Enable and configure Anonymous Blocking in the profile that is used for the Traffic Router.
3. Enable Anonymous Blocking on the Delivery Service.

Upload an Anonymous IP Database

Perform the following steps to upload an anonymous IP database:

Procedure

Step 1 Choose **Provisioning** > **General Settings**.

Step 2 The Anonymous IP Database window appears. If an anonymous IP database is already uploaded, it will be listed in this window. If you need to upload a database, click **here**, click **Browse**, and then browse for the anonymous IP database file.

Note A license is required to use the MaxMind Anonymous IP database. For more information on obtaining this license, please contact your Cisco Account team.

Enable and Configure Anonymous Blocking in the Profile

Perform the following steps to enable and configure anonymous blocking on a Traffic Router profile:

Procedure

Step 1 Choose **Provisioning** > **Edit CDN**. From the **CDN Name** drop-down list, choose the CDN that contains the Traffic Router profiles for which you want to configure Proximity routing.

Step 2 Click the **Profiles** tab and from the Profiles Detail section, click **Edit** for the Traffic Router profile for which you are configuring anonymous blocking.

Note All of the Traffic Routers in a CDN must be assigned the same Traffic Router profile.

Step 3 From the Routing Settings pane that appears, click the **Anonymous Blocking** tab.

Step 4 From the **Anonymous Blocking** tab, use the following table to configure the anonymous blocking parameters. When finished entering the parameters, click **Save** to save the changes.

Table 2: Anonymous Blocking Parameters

Parameter Name	Settings	Default Value	Description
Enable Anonymous Blocking	Enabled (checked) or Disabled (unchecked)	Disabled	This parameter determines whether Anonymous Blocking is enabled or disabled. To use any features of Anonymous Blocking, including the IP4 and IP6 White Lists, you must enable this setting.
Name	Not used by Traffic Router at this time	N/A	This field can be used as a description field.
Customer	Not used by Traffic Router at this time	N/A	This field can be used as a description field.

Parameter Name	Settings	Default Value	Description
Version	N/A	N/A	This is a non-editable field. It keeps track of the number of times the Anonymous Blocking parameters have been modified.
Date	N/A	N/A	This is a non-editable field and it is not used by Traffic Router at this time.
Block Anonymous VPN	Enabled (checked) or Disabled (unchecked)	Disabled	When this option is enabled, the Traffic Router blocks requests from IP addresses that are coming from anonymous VPNs, as identified by the anonymous IP database.
Block Hosting Providers	Enabled (checked) or Disabled (unchecked)	Disabled	When this option is enabled, the Traffic Router blocks requests from IP addresses that are coming from hosting providers, as identified by the anonymous IP database.
Block Public Proxy	Enabled (checked) or Disabled (unchecked)	Disabled	When this option is enabled, the Traffic Router blocks requests from IP addresses that are coming from public proxies, as identified by the anonymous IP database.
Block TorExit Node	Enabled (checked) or Disabled (unchecked)	Disabled	When this option is enabled, the Traffic Router blocks requests from IP addresses that are coming from TorExit nodes, as identified by the anonymous IP database.
IP4 White List	Enter the IPv4 addresses and subnets that you want to ensure Anonymous Blocking does not block.	None	If Anonymous Blocking is enabled, requests from IPv4 addresses that match an entry in the IP White List will be allowed.

Parameter Name	Settings	Default Value	Description
IP6 White List	Enter the IPv6 addresses and prefixes that you want to ensure Anonymous Blocking does not block.	None	If Anonymous Blocking is enabled, requests from IPv6 addresses that match an entry in the IP White List will be allowed.
Polling Interval	Enter the polling interval in milliseconds.	86400000 (1 day)	This setting determines how often the Traffic Router should poll for changes to the Anonymous Blocking configuration.
Redirect URL	Enter the URL to which clients that are blocked by Anonymous blocking should be redirected.	None	If this field is left blank, clients that were blocked by Anonymous blocking will receive a 403 Forbidden message.

Example

The following shows an example of configuring the anonymous blocking settings on the Traffic Router profile.

The screenshot displays the 'Profiles' tab in the Traffic Router configuration interface. Under 'Routing Settings - CCR_CDN_CDN3', the 'Anonymous Blocking' sub-tab is active. The status is 'success'. The 'Enable Anonymous Blocking' checkbox is checked. The 'Name' field is 'OMD-Blocking' and the 'Customer' field is 'CompanyX'. The 'Version' is 0 and the 'Date' is 2017-09-08 18:20:19.476819. The 'Block Anonymous VPN' checkbox is checked, while 'Block Hosting Provider', 'Block Public Proxy', and 'Block TorExit Node' are unchecked. The 'IP4 White List' contains two entries: '10.1.0.0/16' and '192.168.0.0/16'. The 'IP6 White List' contains one entry: '2001:0:0:1::/64'. The 'Polling Interval' is set to '86400000' with a unit of 'ms'. 'Cancel' and 'Save' buttons are at the bottom right.

Enable Anonymous Blocking on a Delivery Service



Note Anonymous Blocking is not supported by DNS Delivery Services unless Edge Geo Blocking is configured. For more information on Edge Geo Blocking, see [Edge Geo Blocking, on page 26](#).

For a Delivery Service to use these anonymous blocking settings, the Delivery Service must be enabled to use anonymous blocking. Perform the following steps to enable anonymous blocking on the Delivery Service:

Procedure

- Step 1** Choose **Provisioning > Edit CDN**. From the CDN Name drop-down list above the tabs, make sure the correct CDN is selected.
- Step 2** Click the **Delivery Services** tab.
- Step 3** From the Delivery Service Name list, choose the Delivery Service for which you want to enable anonymous blocking.
- Step 4** Click the **Advanced Settings** link.
- Step 5** In the Geo Limit panel of the Additional Info window, choose **Yes** from the Anonymous Blocking Enabled drop-down list.
- Step 6** In the Additional Info window click **Save** and then click **Save** at the bottom of the Delivery Services tab.

ASN Blocking



Note ASN Blocking is not supported by DNS Delivery Services unless Edge Geo Blocking is configured. For more information on Edge Geo Blocking, see [Edge Geo Blocking, on page 26](#).

ASN blocking enables you to block client requests based on the autonomous system number (AS number) to which the client IP address belongs. ASN blocking uses the MaxMind ISP or the GeoLite “ASN” database to determine the AS number of an IP address.



Note A license is required to use the MaxMind ISP database. For more information on obtaining this license, please contact your Cisco Account team.

When the profile for the Traffic Router is configured for ASN Blocking, when traffic is received for a Delivery Service that is not blocked by the CZF file, National geoblocking, or Anonymous Blocking, the Traffic Router does the following:

1. Checks whether the Delivery Service has ASN Blocking configured.

2. If the Delivery Service has ASN Blocking configured, the Traffic Router uses the ISP database that has been defined in the Traffic Router profile to determine the AS number to which the requesting client IP address belongs.
3. The Traffic Router checks the ASN of the request against the ASN whitelist configured on the Delivery Service. If the ASN matches an entry in the ASN whitelist, the request is allowed, otherwise the request is blocked and the client is redirected to the Geo Limit Redirect URL if it is configured on the Delivery Service. If the Geo Limit Redirect URL is not configured, the client receives a 403 Forbidden message.

Configure ASN Blocking



Note

ASN Blocking is not supported by DNS Delivery Services unless Edge Geo Blocking is configured. For more information on Edge Geo Blocking, see [Edge Geo Blocking](#).

To configure ASN Blocking, you must do the following:

1. Upload an ISP database to the CDN that will be used to determine the AS number to which the requesting client IP address belongs.
2. Configure the Traffic Router profile to use an ISP database that has been uploaded to the CDN.
3. Configure the ASN whitelist on the Delivery Service to enable ASN blocking and to identify the AS numbers from which client requests should be allowed. If the ASN whitelist is not configured, ASN blocking will not block any client requests.

Upload an ISP Database

Perform the following steps to upload an ISP database to the CDN:

Procedure

- Step 1** Copy the ISP database that you received from MaxMind to the system from which you are running the OMD Director GUI.

Note A license is required to use the MaxMind ISP database. For more information on obtaining this license, please contact your Cisco Account team.
- Step 2** Choose **Provisioning > General Settings > ISP Database**.
- Step 3** The ISP Database window appears. If an ISP database is already uploaded, it will be listed in this window.
- Step 4** If you need to upload an ISP database, click the **Add (+)** icon. In the ISP Database dialog box that appears, click **Browse** and then browse for the ISP database file.
- Step 5** Click **Save** to finish adding the ISP database to the CDN.

Assign an ISP Database to the Traffic Router Profile

Perform the following steps to assign an ISP database to the Traffic Router profile:

Procedure

-
- Step 1** Choose **Provisioning** > **Edit CDN**. From the CDN Name drop-down list, choose the CDN that contains the Traffic Router profile for which you are configuring ASN blocking.
- Step 2** Click the **Profiles** tab and from the Profiles Detail section, click **Edit** for the Traffic Router profile for which you are configuring ASN blocking.
- Note** All of the Traffic Routers in a CDN must be assigned the same Traffic Router profile.
- Step 3** From the Routing Settings pane that appears, click the **ISP Database** tab.
- Step 4** From the ISP Database drop-down list, choose the ISP database to assign to the Traffic Router profile.
- Step 5** Click **Save** to save the changes.
-

Enable ASN Blocking on a Delivery Service



Note ASN Blocking is not supported by DNS Delivery Services unless Edge Geo Blocking is configured. For more information on Edge Geo Blocking, see [Edge Geo Blocking, on page 26](#).

For a Delivery Service to use the ISP database to perform ASN blocking, the Delivery Service must be enabled to use ASN blocking. To enable ASN blocking on a Delivery Service, you must configure the ASN whitelist of the Delivery Service.

Perform the following steps to configure the ASN whitelist, thereby enabling ASN blocking on the Delivery Service:

Procedure

-
- Step 1** Choose **Provisioning** > **Edit CDN**. From the CDN Name drop-down list above the tabs, make sure the correct CDN is selected.
- Step 2** Click the **Delivery Services** tab.
- Step 3** From the Delivery Service Name list, choose the Delivery Service for which you want to enable ASN blocking.
- Step 4** Click the **Advanced Settings** link.
- Step 5** In the Geo Limit panel of the Additional Info window, in the ASN Whitelist field, enter a comma separated list of AS numbers from which the client requests should be allowed.
- Step 6** In the Additional Info window click **Save** and then click **Save** at the bottom of the Delivery Services tab.
- Note** If you leave the ASN Whitelist empty, ASN blocking is *not* enabled and no client requests will be blocked based on their AS number.
-

NGB Whitelist

When the Geo Limit setting of a Delivery Service is set to "CZF + CountryCode(s)", this enables National geoblocking (NGB). When NGB is enabled, if the client IP address does not match an entry in the CZF file, the Traffic Router uses a geolocation database to determine the country of the requesting client IP address. If the client IP address is from a country that matches an entry in the "Geo Limit Country Codes" setting of the Delivery Service, the Traffic Router allows the request.

When NGB is enabled, by default if the client IP address is not in the CZF file and the client IP address is *not* from a country listed in the "Geo Limit Country Codes" field, the Traffic Router will reject the client request. However, you can configure an NGB whitelist to allow a client request that NGB would otherwise block.



Note

When the Geo Miss Default Latitude and Longitude settings of the Delivery Service are configured, these settings are used to determine the country for any client IP addresses that are *not* recognized by the geolocation database. Therefore, the NGB whitelist is typically used to allow IPs that *are* recognized by geolocation but are *not* in the allowed Country Codes as listed in the "Geo Limit Country Codes" field.

For information on configuring the NGB whitelist, see [Configure an NGB Whitelist](#).

Configure an NGB Whitelist

Perform the following steps to create an NGB whitelist and assign it to the Traffic Router profile:

Procedure

- Step 1** On a system from which you can access the OMD Director GUI, create an NGB whitelist JSON file. For the syntax of the NGB whitelist JSON file and for an example of this file, see [NGB Whitelist File, on page 265](#).
 - Step 2** From the system on which the NGB whitelist JSON file was created, log in to the OMD Director GUI as an administrator or CDN Admin.
 - Step 3** Choose **Provisioning** > **Edit CDN**. From the CDN Name drop-down list, choose the CDN that contains the Traffic Router profile for which you want to upload an NGB whitelist.
 - Step 4** Click the **Profiles** tab and from the Profiles Detail section, click **Edit** for the Traffic Router profile for which you are configuring the NGB whitelist.
- Note** All of the Traffic Routers in a CDN must be assigned the same Traffic Router profile.
- Step 5** From the Routing Settings pane that appears, click the **NGB Whitelist** tab.
 - Step 6** Click the **Upload** link.
 - Step 7** In the NGB Whitelist File Upload field, click **Browse**, browse for the NGB whitelist JSON file that was created in Step 1, and then click **Open**.
 - Step 8** Click **Save** to save the changes.



CHAPTER 9

Manage Content Invalidation Policies

After an object is cached, it normally remains in the cache until it expires or is evicted to make room for new content. Sometimes you may need to make content unavailable to the client before it expires, such as if content is removed from an Origin Server or a VOD asset has expired. By creating content invalidation policies in Media Streamer, you can configure the CDN to selectively make cached content inaccessible to the client without the overhead of removing the content from all of the cache servers.



Note To support this feature, the Origin Server must support HTTP/1.1 cache revalidate and invalidated content must be removed from the origin sites.



Note If you are running OMD Director in an HA configuration, you can only make configuration changes from the Primary OMD Director instance. To perform the steps in this section, make sure you are logging into the Primary OMD Director instance.

This chapter describes how to manage the content invalidation policies and includes the following topics:

- [Add a Content Invalidation Policy, on page 177](#)
- [Delete a Content Invalidation Policy, on page 179](#)

Add a Content Invalidation Policy

Perform the following steps to create a Content Invalidation policy:

Procedure

- Step 1** Choose **Provisioning > Edit CDN**.
- Step 2** Make sure the correct CDN is selected in the **CDN Name** drop-down list and click the **Content Invalidation** tab.
- Step 3** A window appears that shows the current Content Invalidation policies. This window displays the Delivery Service to which the content invalidation policies applies, what content the policy applies to, when the policy begins, and how long the policy will be in effect.

Add a Content Invalidation Policy

Provisioning > Edit CDN

CDN Name:

Servers Cache Groups Delivery Services Client Routing Profiles **Content Invalidation**

Invalidation Content

Show 10 entries

Delivery Service Name	Regex	Start Time	TTL	
vod	http://10.74.25.168/video/May1/*	2017-12-12 22:59:59+00	54h	<input type="button" value="+"/>
vod1	http://10.74.25.168/video/May1/*	2017-12-12 22:59:59+00	54h	<input type="button" value="⌵"/>
vod2	http://10.74.25.168/video/May1/* /.*	2017-11-10 07:34:21+00	48h	<input type="button" value="⌵"/>

Showing 1 to 3 of 3 entries

Previous 1 Next

Step 4 To create a new content invalidation policy, click the Create New Content Invalidation (+) icon in the header of the last column.

Step 5 The Add Invalidation Content window appears. In this window configure the following information:

- **Delivery Service Name:** Choose the Delivery Service to which this content invalidation policy applies.
- **Path Regex:** Enter the path to the content that should be invalidated for this Delivery Service. Do not enter the Origin Server Base URL, only enter the path to the content, for example /video/dec10/*. The Origin Server Base URL that is configured for the Delivery Service is automatically added to create the complete Regex value. This field enables you to choose which content for the Delivery Service should be invalidated so you do not have to invalidate all of the content for the Delivery Service.
- **Time to Live (in hours):** Enter the number of hours the content validation rule should be activate. The default value is 54 hours.
- **Start Time:** Use the calendar and time picker to choose the date and time when this policy should begin. The policy will stay in effect until the number of hours specified in the Time to Live field have passed. For example, if the Time to Live is 72 hours and the Start Time is 1:00 am on January 1, 2018, the policy will expire on 1:00 am on January 4, 2018.

Note The Start Time must be within the next 2 days.

Step 6 When you are done entering the values for the new content invalidation policy, click **Add**. The following is an example:

Add Invalidation Content

Delivery Service Name *

Path Regex *

Time To Live (in hours) *

Start Time *

Note For information on using the Cisco OMD Director REST APIs to add a content invalidation policy, see [Manage Content Invalidation using the OMD Director REST API, on page 339](#).

Delete a Content Invalidation Policy

Perform the following steps to delete a Content Invalidation policy:

Procedure

- Step 1** Choose **Provisioning > Edit CDN**.
- Step 2** Make sure the correct CDN is selected in the **CDN Name** drop-down list and click the **Content Invalidation** tab.
- Step 3** In the last column, click the **Delete** (trashcan) icon for the content invalidation policy that you would like to delete.
- Step 4** In the Deleting Content Invalidation confirmation window that appears, click **Yes** to confirm that you want to delete the policy.

Note For information on using the Cisco OMD Director REST APIs to delete a content invalidation policy, see [Manage Content Invalidation using the OMD Director REST API, on page 339](#).



CHAPTER 10

Device Groups

Device groups enable you to group servers together to more easily assign them to a Delivery Service. A device group can contain Edge cache servers, Mid cache servers, and Origin Servers and can include servers from more than one Media Streamer CDN. When a device group is assigned to a Delivery Service, only the servers in that device group that belong to the same CDN as the Delivery Service will serve the Delivery Service.

This chapter describes device groups and how to manage them, and includes the following topics:

- [Device Groups Overview, on page 181](#)
- [Create a Device Group, on page 182](#)
- [Edit a Device Group, on page 184](#)
- [Delete a Device Group, on page 185](#)

Device Groups Overview



Note

If you are running OMD Director in an HA configuration, you can only make configuration changes from the Primary OMD Director instance. To perform the steps in this section, make sure you are logging into the Primary OMD Director instance.

A device group is different from a cache group. A cache group is a logical grouping of servers used to provide high availability and to provide a hierarchy of cache servers. Cache groups are also assigned geographical coordinates that can be used to determine the best cache group to deliver the content. The main purpose of a device group is to group servers together so you can more easily assign the servers to a Delivery Service.

Device groups do not replace cache groups; device groups are an additional tool to help more efficiently assign servers to Delivery Services. Device groups and cache groups have the following differences:

- A server can only belong to one cache group but it can belong to multiple device groups.
- A server has to belong to a cache group but it does not have to belong to a device group.
- A cache group can only contain servers of the same type, while a device group can contain different types of servers.
- A device group can contain servers from more than one CDN. A cache group can only contain servers from the same CDN.

- An Edge cache group is assigned a Mid cache group as a parent and secondary parent to create a hierarchy. Device groups do not have any parents assigned.
- The geographical coordinates of the cache group determine the location of the servers in that cache group. Device groups are not assigned geographical coordinates.

If more than one device group is applied to the Delivery Service, the servers from both device groups will serve the Delivery Service. If specific servers are assigned in the Assign Cache Servers field and device groups are assigned to a Delivery Service, then the union of the specific servers and the servers in the device groups will all service the Delivery Service.

Create a Device Group

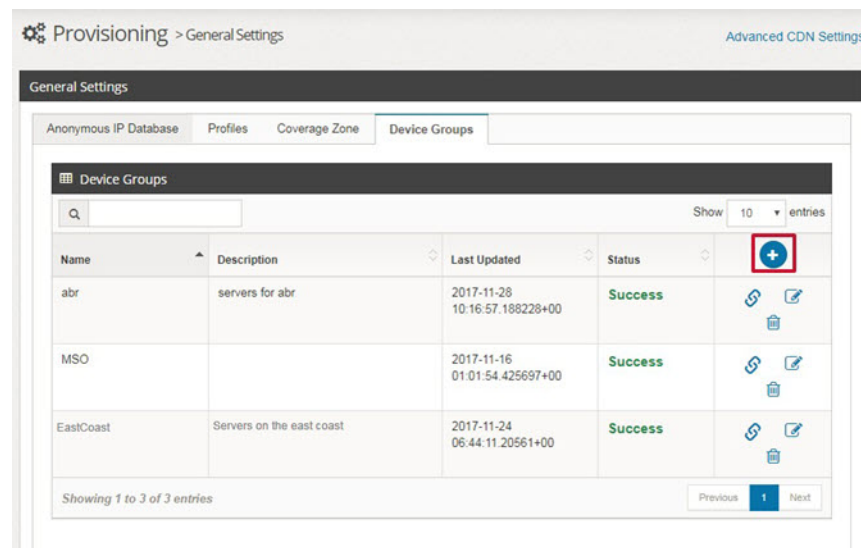
To create a device group and assign servers to the group, perform the following steps:

Procedure

Step 1 Choose **Provisioning > General Settings**.

Step 2 Click the **Device Groups** tab. The Device Groups tab displays a list of device groups, including a description of the group, when it was last updated, and its status.

Step 3 To add a device group, click the Add (+) icon.



Step 4 The Add Device Group window appears. In this window enter a name and description for the device group.

Step 5 Click **Add** to add the device group.



Add Device Group

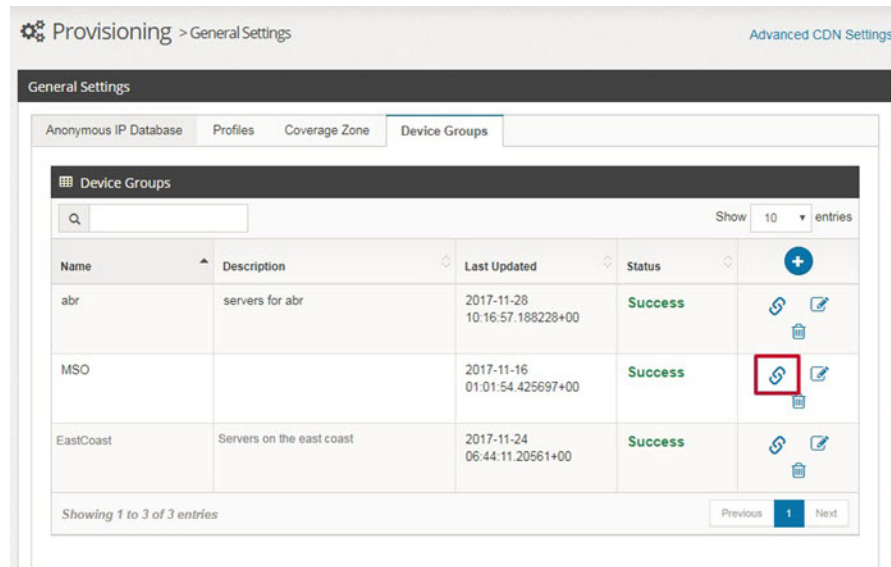
Name *
WestCoast

Description *
Servers on the west coast

Close Add

Step 6

To add servers to the device group, in the right most column of the device group, click the **Assign Servers** icon.






Provisioning > General Settings

General Settings

Anonymous IP Database Profiles Coverage Zone Device Groups

Device Groups

Show 10 entries

Name	Description	Last Updated	Status	
abr	servers for abr	2017-11-28 10:16:57.188228+00	Success	
MSO		2017-11-16 01:01:54.425697+00	Success	
EastCoast	Servers on the east coast	2017-11-24 06:44:11.20561+00	Success	

Showing 1 to 3 of 3 entries

Previous 1 Next

Step 7

The Assign Servers to Device Group window appears. From this window in the Assign Servers to Device Group section, configure the following:

- **Assign Cache Servers:** From this drop-down list, choose a server to assign to the device group.
- **Assign Cache Servers IP:** If the server you are assigning is an Edge cache server, choose the IP address that you would like the Edge cache server to use when streaming content for Delivery Service to which this device group is assigned. By default only the primary IP address is available. For information on how to add a secondary IP address for the Edge cache server to use for streaming content, see [Manage Secondary Streaming IPs, on page 68](#).

Note Mid cache servers and origin servers can only use their primary IP address.

Step 8

To add additional servers to the device group, click the + (Add) icon and repeat Step 7.

Step 9

When you are done adding the servers to the group, click **Save**.

- Step 10** When the device group has been successfully updated, the “Processing” status in the title bar will disappear and you will see a green check mark for every action in the Status column. Click **OK** to close the status window.

Edit a Device Group

Procedure

- Step 1** To change the name or description of a device group, perform the following steps:
- Choose **Provisioning > General Settings** and click the **Device Group** tab.
 - From the Device Groups list, in the right most column of the device group, click the **Edit** icon for the server you want to edit.
 - In the Edit Device Group box make the necessary changes and click **Save**.
- Step 2** To change which servers are assigned to the device group, perform the following steps:
- Click the **Assign Servers** icon.
 - To change the IP address that an Edge cache server uses to stream content for Delivery Services to which this device group is assigned, choose the IP address from the Assign Cache Servers IP drop-down list. By default only the primary IP address is available. For information on how to add a secondary IP address for the Edge cache server to use for streaming content, see [Manage Secondary Streaming IPs, on page 68](#).
- Note** Mid cache servers and origin servers can only use their primary IP address.
- To remove a server from the device group, in the Assign Servers to Device Group section, click the - (Delete) icon at the end of the row for the server you want to delete.
 - To add a server to the device group, in the Assign Servers to Device Group section, click the + (Add) icon at the end of the row and configure the following:
 - Assign Cache Servers:** From this drop-down list, choose a server to assign to the device group.
 - Assign Cache Servers IP:** If the server you are assigning is an Edge cache server, choose the IP address that you would like the Edge cache server to use when streaming content for Delivery Services to which this device group is assigned. By default only the primary IP address is available. For

information on how to add a secondary IP address for the Edge cache server to use for streaming content, see [Manage Secondary Streaming IPs, on page 68](#).

Note Mid cache servers and origin servers can only use their primary IP address.

- e) Click **Save** to save the changes.
 - f) When the device group has been successfully updated, the “Processing” status in the title bar will disappear and you will see a green check mark for every action in the Status column. Click **OK** to close the status window.
-

Delete a Device Group

To delete a device group, perform the following steps:

Procedure

- Step 1** Choose **Provisioning > General Settings** and click the **Device Group** tab.
 - Step 2** From the Device Groups list, in the right most column of the device group, click the **Delete** icon for the server you want to delete.
 - Step 3** In the Deleting Device Group confirmation window that appears, click **Yes** to confirm that you want to delete the device group.
-



CHAPTER 11

KPI Metrics

Using graphs and statistics, OMD Director enables you to monitor your CDN caching performance. The KPI Metrics menu in OMD Director enables you to view the key analytics information of all of the CDN Edge caches and Mid caches.

This chapter looks at the different information that you can view from the KPI Metrics menu in OMD Director and includes the following topics:

- [Getting Started, on page 187](#)
- [KPI Metrics > CDN Edge, on page 188](#)
- [KPI Metrics > CDN Mid, on page 195](#)

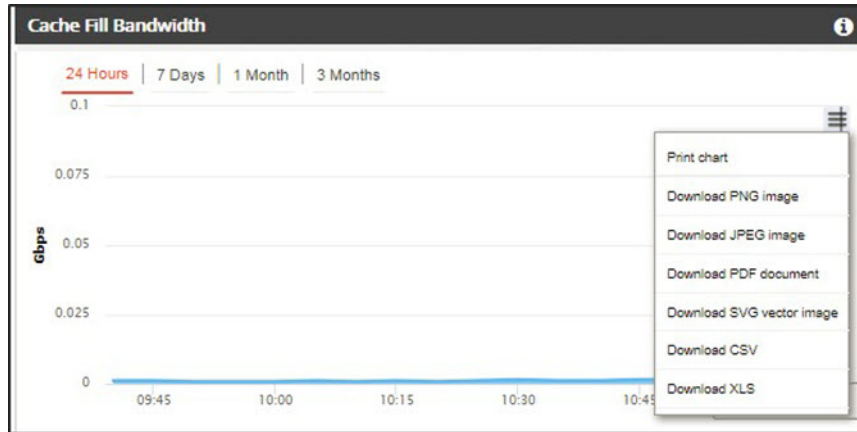
Getting Started

The graphs that are available from the KPI Metrics menus include the following interactive graph elements:

- **Time Frame:** You can choose the time frame across which you want to display statistics in the graph by clicking the time frame at the top of the graph.
 - **24 hour:** When you select this time range, the graph displays the last 24 hours of data. To refresh the data, click the Refresh icon in the upper right corner.
 - **7 days:** When you select this time range, the graph displays the previous 7 days of data and the graph is updated once a day.
 - **1 Month:** When you select this time range, the graph displays the previous month of data and the graph is updated once a day.
 - **3 Months:** When you select this time range, the graph displays the previous three months of data and the graph is updated once a day.
- By hovering over the Expand icon in the upper-right of the graph you can do the following:
 - **Print:** You can print the graph.
 - **Download the graph:** You can download the graph as any of the following file types:
 - PNG
 - JPEG
 - PDF

- SVG
- CSV
- XLS

- **Mouse Over Time Value:** If you hover the mouse over the graph at a specific point, it will show you the value represented in the graph for that specific date and time. The following is an example of the graphing options available:



Note If there is no data for any of the graphs, they will show “No data to display”.

The KPI Metrics menu contains the following options:

- CDN Edge
- CDN Mid

The following sections describe the information that is available from each of these menu options.

KPI Metrics > CDN Edge

By default, the KPI Metrics > CDN Edge page provides a topology of the CDN Edge cache groups and provides statistical information across all of the Edge cache groups in the CDN for ABR traffic. The CDN Edge Cache contains the following graphs:



Note If there is no data for any of the graphs, they will show “No data to display”.

CDN Topology

The CDN Topology pane displays the geographic topology of the CDN Edge cache groups. The dots on the map represent the cache groups based on their configured latitude and longitude. The color of the dots represent the Cache Hit Ratio (CHR) percentage:

- Green: CHR is \geq to 70%
- Yellow: CHR is \geq 40% and $<$ 70%
- Red: CHR is $<$ 40%



From this pane you can do the following:

- Double-click the map to zoom into a specific region
- Move the focus of the map by clicking and dragging the map
- Use the Zoom In and Zoom Out tools to zoom in and out
- Click the Home button to return the map to its original view

You can also hover over a specific Edge cache group icon to see the following information about the group:

- Name of the group
- Average Cache Hit Ratio (CHR) of the group
- Average delivery bandwidth
- Cache servers in the group and their individual CHR and delivery bandwidth averages

**Note**

If the transaction logs configuration is incorrect or logs are not generated on the server, then data analytics for that server will not be available. In this case, the servers will not appear when you hover over the cache group icon.

Delivery Bandwidth

The speedometer in the Delivery Bandwidth panel displays the average bandwidth over the last minute from Edge cache servers to clients.

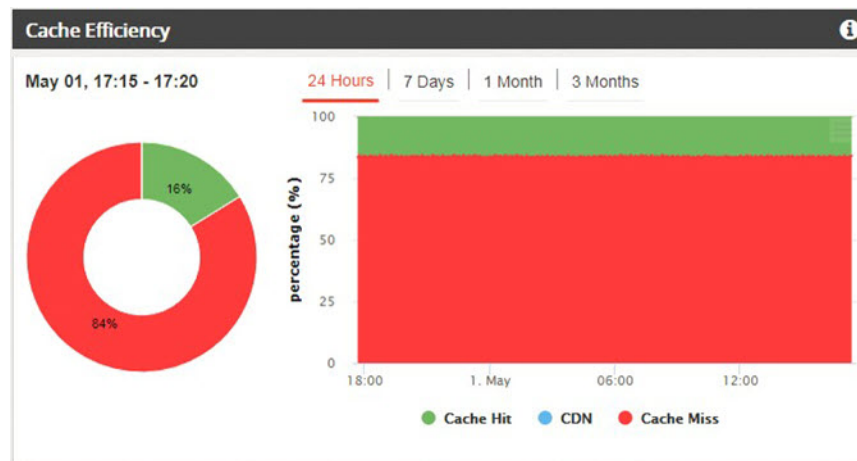
The graph displays the same information across the selected time range with the ability to drill down to a specific date and time and contains all of the interactive graph elements that were discussed in [Getting Started, on page 187](#).



Cache Efficiency

The donut chart on the Cache Efficiency panel displays the ratio of Edge cache hits and misses, averaged over the last 5 minutes.

The graph displays the same information across the selected time range with the ability to drill down to a specific date and time. This graph contains all of the interactive graph elements that were discussed in [Getting Started, on page 187](#).



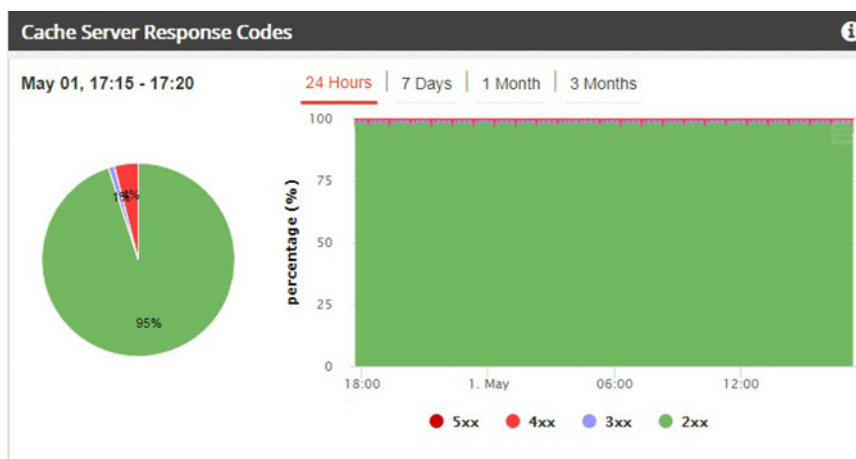
Cache Server Response Codes

The pie chart on the Cache Server Response Codes panel displays the percentage of the following HTTP response codes received by the clients from the Edge caches over the last 5 minutes:

- **2xx (Success) codes:** These codes indicate that the client request was successfully received, understood, and accepted.
- **3xx (Redirection) codes:** These codes indicate that the user agent (a web browser or a crawler) needs to take further action when trying to access a particular resource. Generally the user agent is automatically forwarded or redirected to another resource (URL) without interaction of the user.
- **4xx (Client Error) codes:** These error codes indicate that content was not found on the cache. This can occur for two main reasons:
 - The URL requested by the client is not available on the Origin Server, for example, the manifest URL is not correct.
 - The cache is missing the Delivery Service configuration.
- **5xx (Server Error) codes:** These error codes generally indicate an internal error either in the CDN or on the Origin server.

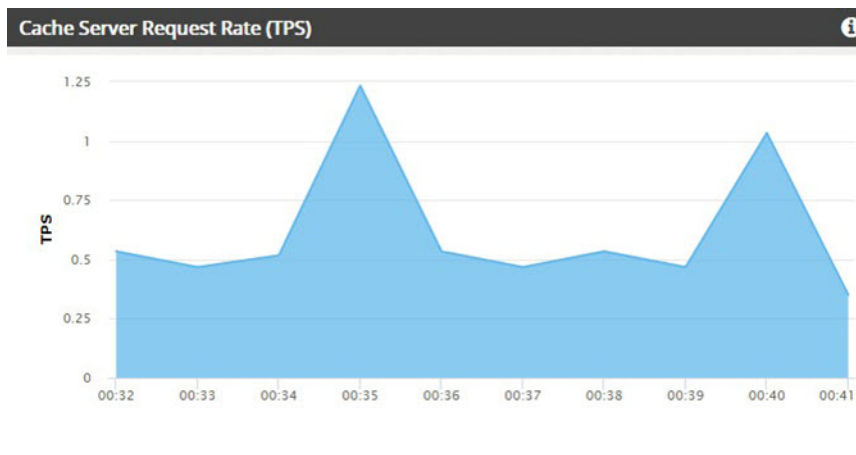
For both 4xx and 5xx errors, you can use Insights reports to identify the set of requests that are generating 4xx and 5xx errors. After you have the details of those requests, you can use that information to troubleshoot the Origin Server and CDN.

The graph shows the same information across the selected time range with the ability to drill down to a specific date and time. This graph contains all of the interactive graph elements that were discussed in [Getting Started](#), on page 187.



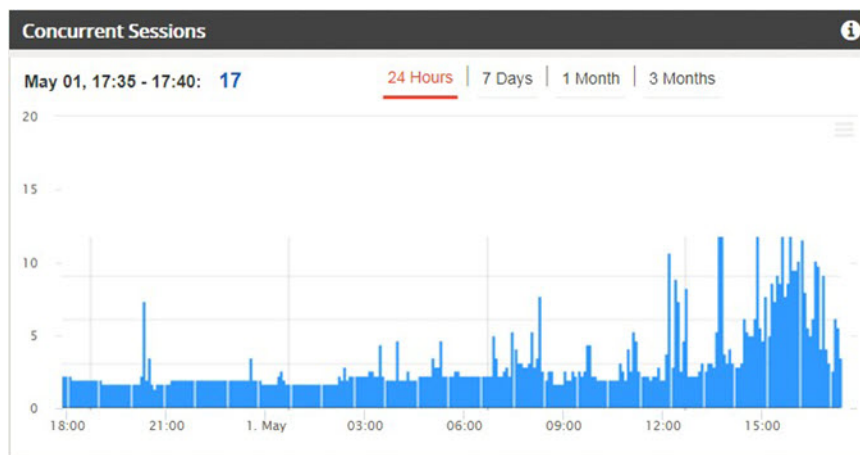
Cache Server Request Rate

The Cache Server Request Rate graph displays the number of Transactions per Second (TPS) between the Edge caches and clients for the last 10 minutes with the ability to drill down to a specific time. This chart updates every minute.



Concurrent Sessions

The Concurrent Sessions graph displays the number of current sessions between the Edge caches and clients with the ability to drill down to a specific date and time. This pane also displays the total number of current sessions that have been seen for the indicated 5 minute range in the top left of the pane. For example, in the graph below, the total number of concurrent sessions from 23:15 to 23:20 on February 14 was 8. If there is no data to display for this value, it will show “N/A”. This graph contains all of the interactive graph elements that were discussed in [Getting Started, on page 187](#).



Unique Clients

The Unique Clients graph displays the number of unique client IP addresses to which the Edge caches have served requests with the ability to drill down to a specific date and time. This pane also displays the total number of unique clients that have been seen for the indicated 5 minute range in the top left of the pane. For example, in the graph below, the total number of unique clients from 23:15 to 23:20 on February 14 was 7. If there is no data to display for this value, it will show “N/A”. This graph contains all of the interactive graph elements that were discussed in [Getting Started, on page 187](#).



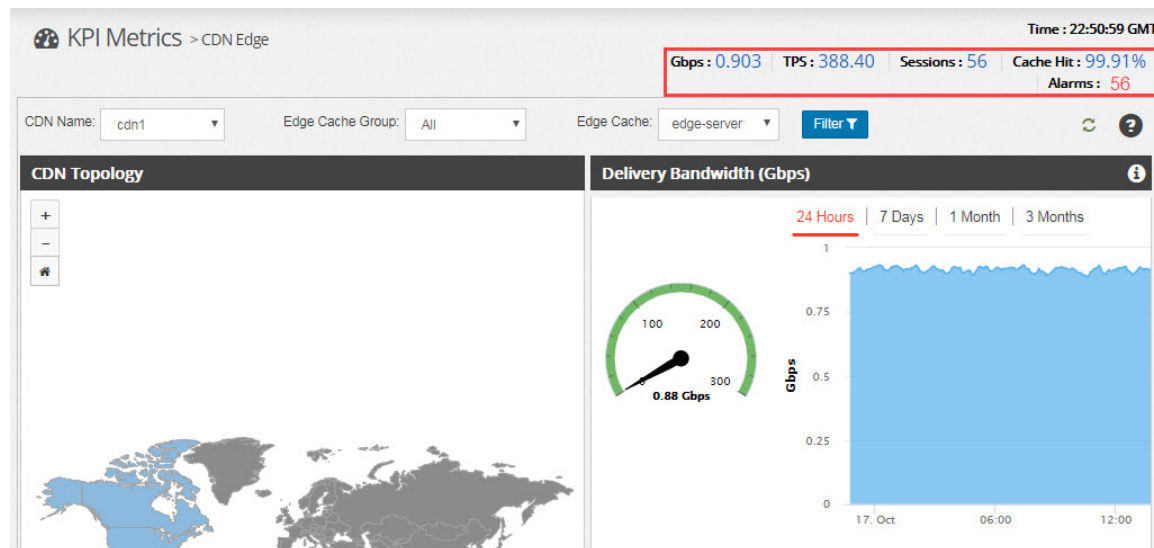
System Overview

The top right of the KPI Metrics > CDN Edge page provides the following generic CDN information across all of the Edge cache servers, based on a summary search over the last 5 minutes. These values do not change when you filter which Edge cache groups or Edge cache servers you are viewing:

- The current delivery bandwidth
- The current transactions per second (TPS)
- The current number of cache sessions

- The current Edge cache hit and miss ratio
- The current number of active alarms

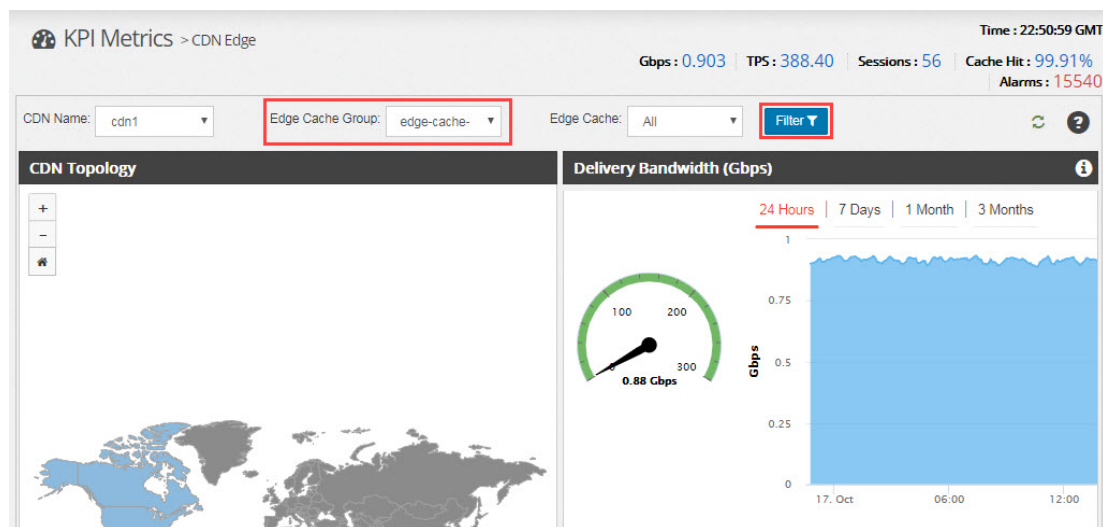
The following is an example:



Filtering the CDN Edge Graphs

If you would like to see the statistical information for one specific Edge cache group or one specific Edge cache server, you can use the filtering options at the top of the KPI Metrics > CDN Edge page:

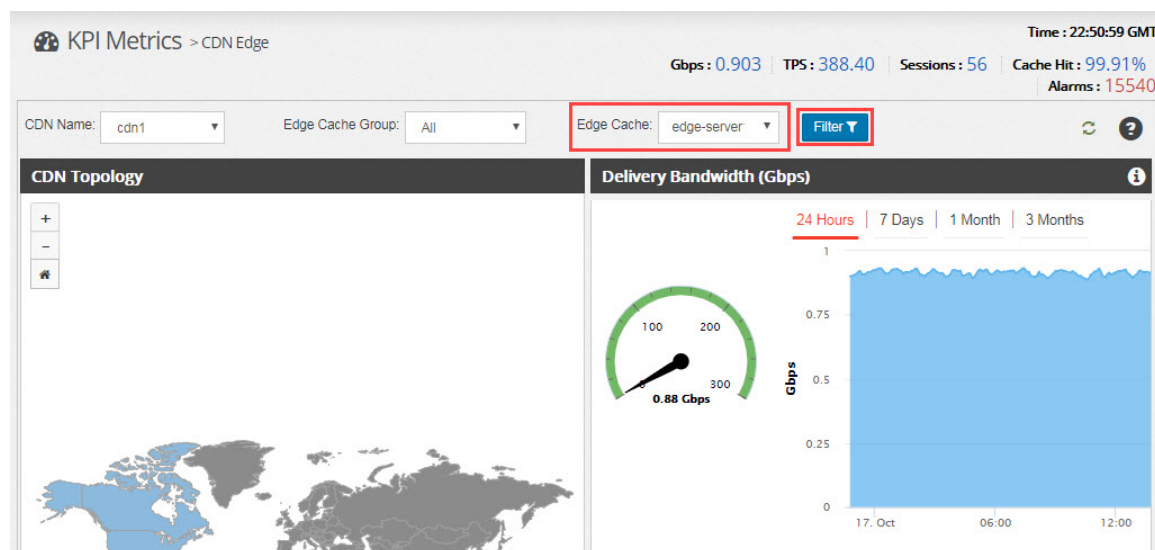
- To filter based on the Edge cache group, from the Edge Cache Group drop-down list, choose the Edge cache group you want to filter on and click **Filter**. All of the graphs on the page are filtered except for the Topology graph. The following is an example:



- To filter based on the Edge cache server, from the Edge Cache drop-down list, choose the Edge cache server you want to filter on and click **Filter**.

- If you choose a group in the Edge Cache Group drop-down list, the Edge Cache list will only show Edge cache servers in that group.
- If you choose All in the Edge Cache Group list, the Edge Cache list will show all of the Edge cache groups and which servers are in each cache group. All of the graphs on the page are filtered except for the Topology graph.

The following is an example:



KPI Metrics > CDN Mid

By default, the KPI Metrics > CDN Mid page provides a topology of the CDN Mid cache groups and provides statistical information across all of the Mid cache groups in the CDN for ABR traffic. The CDN Mid Cache page contains the following graphs:



Note

If there is no data for any of the graphs, they will show “No data to display”.

CDN Topology

The CDN Topology pane displays the geographic topology of the CDN Mid cache groups. The dots on the map represent the cache groups based on their configured latitude and longitude. The color of the dots represent the Cache Hit Ratio (CHR) percentage:

- **Green:** CHR is \geq 70%
- **Yellow:** CHR is \geq 40% and $<$ 70%
- **Red:** CHR is $<$ 40%



From this pane you can do the following:

- Double click the map to zoom into a specific region
- Move the focus of the map by clicking and dragging the map
- Use the Zoom In and Zoom Out tools to zoom in and out
- Click the Home button to return the map to its original view

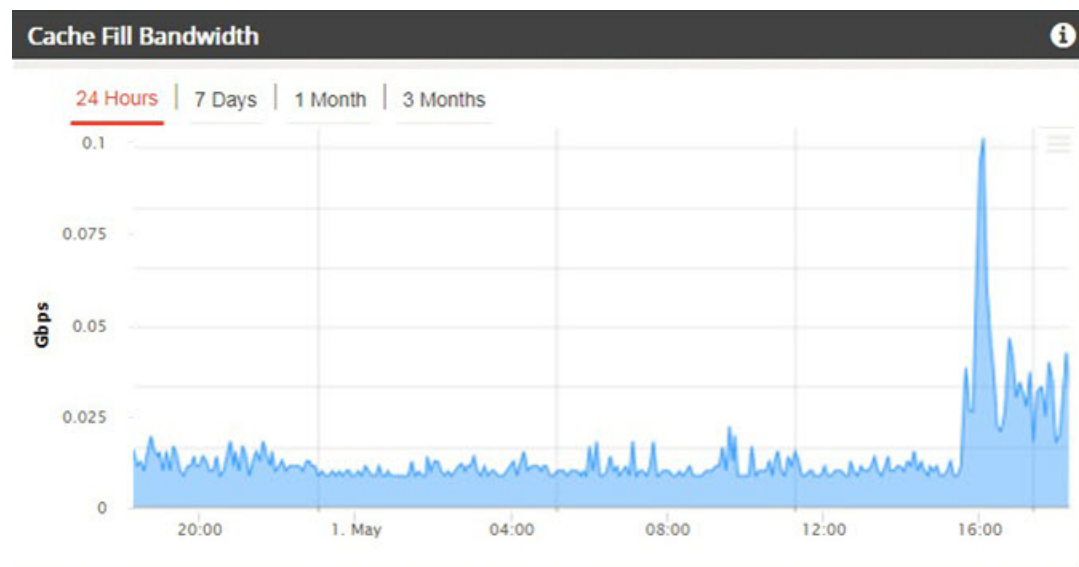
You can also hover over a specific Mid cache group icon to see the following information about the group:

- Name of the group
- Average Cache Hit Ratio (CHR) of the group
- Cache fill bandwidth
- Cache servers in the group and their individual CHR and delivery bandwidth averages

Cache Fill Bandwidth

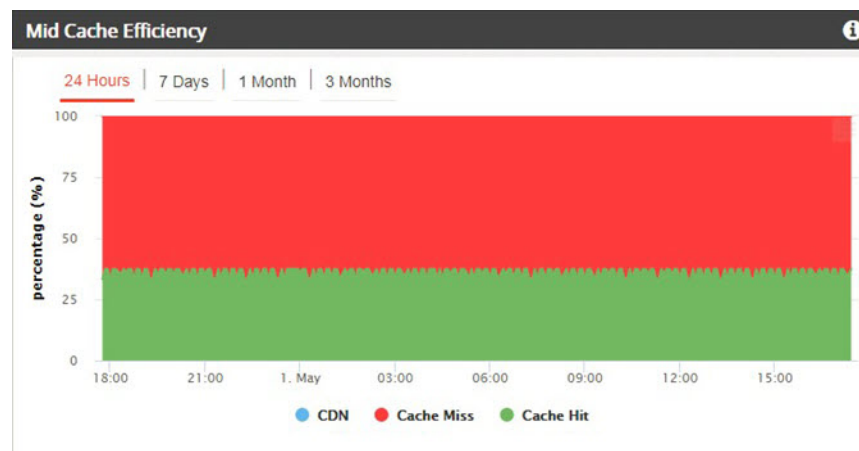
The Cache Fill Bandwidth graph displays the average inbound bandwidth from Mid caches to Edge caches across the selected time range with the ability to drill down to a specific date and time. This graph contains all of the interactive graph elements that were discussed in the "Getting Started" section.

Figure 4: Cache Fill Bandwidth



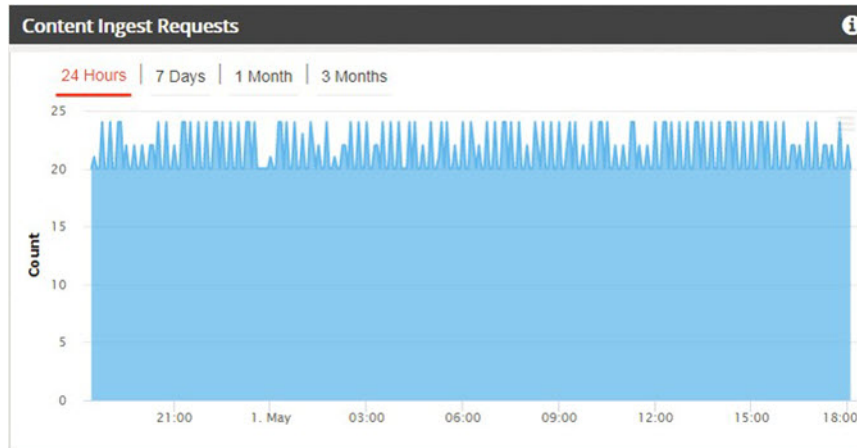
Mid Cache Efficiency

The Mid Cache Efficiency graph displays the ratio of Mid cache hits and misses for cache-fill requests from the edge caches across the selected time range with the ability to drill down to a specific date and time. This graph contains all of the interactive graph elements that were discussed in [Getting Started, on page 187](#).



Content Ingest Requests

The Content Ingest Requests graph displays the number of HTTP objects that the Mid caches have ingested from the Origin Server across the selected time range with the ability to drill down to a specific date and time. This graph contains all of the interactive graph elements that were discussed in [Getting Started, on page 187](#).

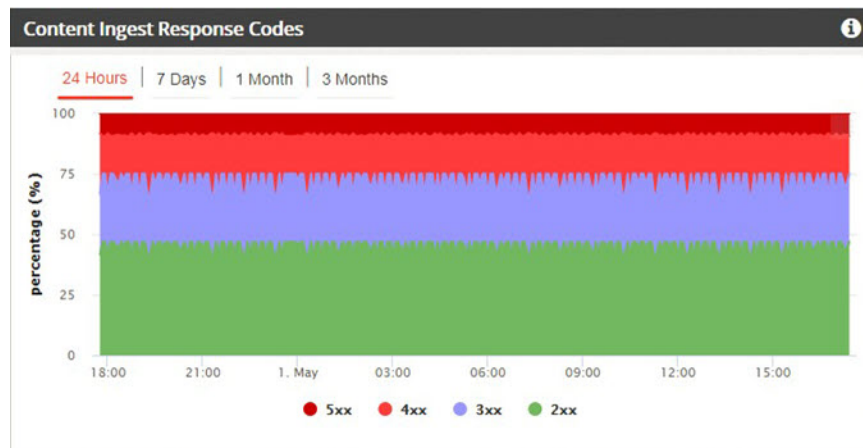


Content Ingest Response Codes

The Content Ingest Response Codes graph displays the percentage of the following HTTP response codes received by the Edge caches from the Mid caches across the selected time range with the ability to drill down to a specific date and time:

- **2xx (Success) codes:** These codes indicate that the client request was successfully received, understood, and accepted.
- **3xx (Redirection) codes:** These codes indicate that the user agent (a web browser or a crawler) needs to take further action when trying to access a particular resource. Generally the user agent is automatically forwarded or redirected to another resource (URL) without interaction of the user.
- **4xx (Client Error) codes:** These error codes indicate that content was not found on the cache. This can occur for two main reasons:
 - The URL requested by the client is not available on the Origin Server, for example, the manifest URL is not correct.
 - The cache is missing the Delivery Service configuration.
- **5xx (Server Error) codes:** These error codes generally indicate an internal error either in the CDN or on the Origin server.

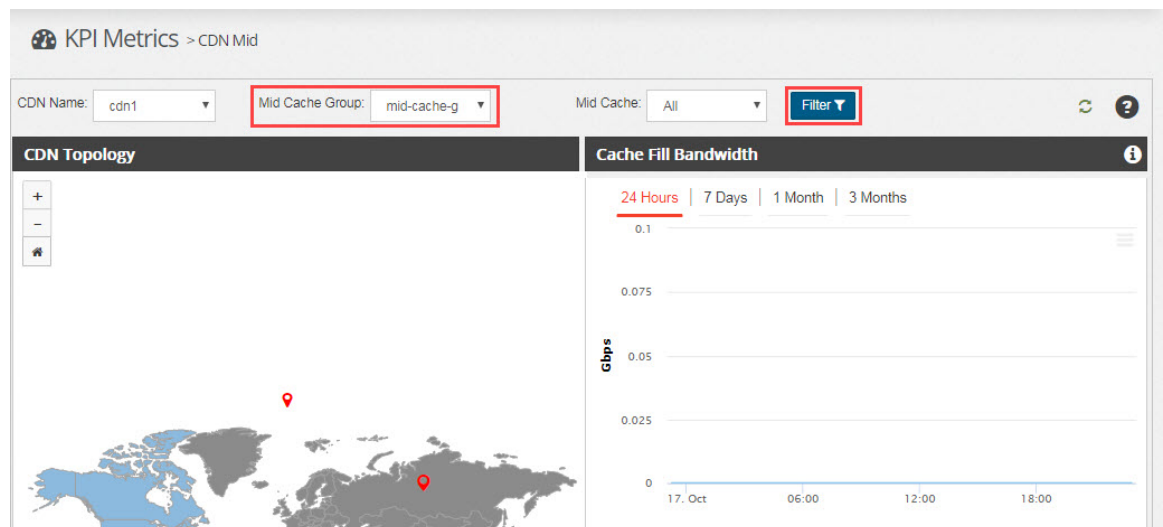
For both 4xx and 5xx errors, you can use Insights reports to identify the set of requests that are generating 4xx and 5xx errors. After you have the details of those requests, you can use that information to troubleshoot the Origin Server and CDN. This graph contains all of the interactive graph elements that were discussed in [Getting Started, on page 187](#).



Filtering the CDN Mid Graphs

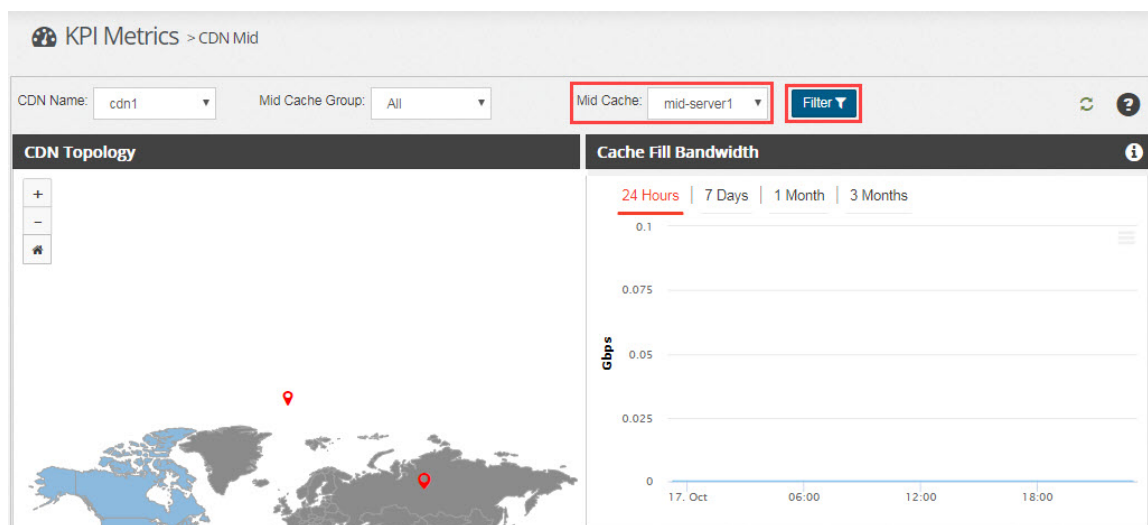
If you would like to see the statistical information for one specific Mid cache group or one specific Mid cache server, you can use the filtering options at the top of the KPI Metrics > CDN Mid page:

- To filter based on the Mid cache group, from the **Mid Cache Group** drop-down list, choose the Mid cache group you want to filter on and click **Filter**. All of the graphs on the page are filtered except for the Topology graph.



- To filter based on the Mid cache server, from the Mid Cache drop-down list, choose the Mid cache server you want to filter on and click **Filter**.
 - If you choose a group in the Mid Cache Group drop-down list, the Mid Cache list will only show Mid cache servers in that group.
 - If you choose All in the Mid Cache Group list, the Mid Cache list will show all of the Mid cache groups and which servers are in each cache group. All of the graphs on the page are filtered except for the Topology graph.

The following is an example:





CHAPTER 12

CDN Monitoring

The Monitoring menu in OMD Director enables you to view system-level statistics for individual cache servers and CDN alarms that have been generated.

This chapter looks at the different information that you can view from the Monitor menu in OMD Director.



Note

If OMD Director is running in an HA configuration, you can view and filter all of the information available in the Server Metrics and Alarms pages from both the Primary and Backup OMD Director instance. However, you cannot acknowledge or clear alarms, configure rules for alarms, or mute alarms if you are logged into an OMD Director instance running in Backup mode. If your OMD Director instance is in Detached mode because of a Director instance failure, you can view and filter all of the information available in the Server Metrics and CDN Alarms pages and you can also acknowledge, clear, and save alarms.

This chapter includes the following topics:

- [Server Metrics, on page 201](#)
- [Alarms, on page 205](#)

Server Metrics

Server metrics enables you to see a server-based view of system-level statistics for a specific cache server. To see these metrics, choose **Monitor > Server Metrics**. When you choose Monitor > Server Metrics, OMD Manager opens a pane that retrieves the server metrics from OMD Monitor.

The graph displays the same information across the selected time range with the ability to drill down to a specific date and time. This graph contains all of the interactive graph elements that were discussed in [Getting Started, on page 187](#).



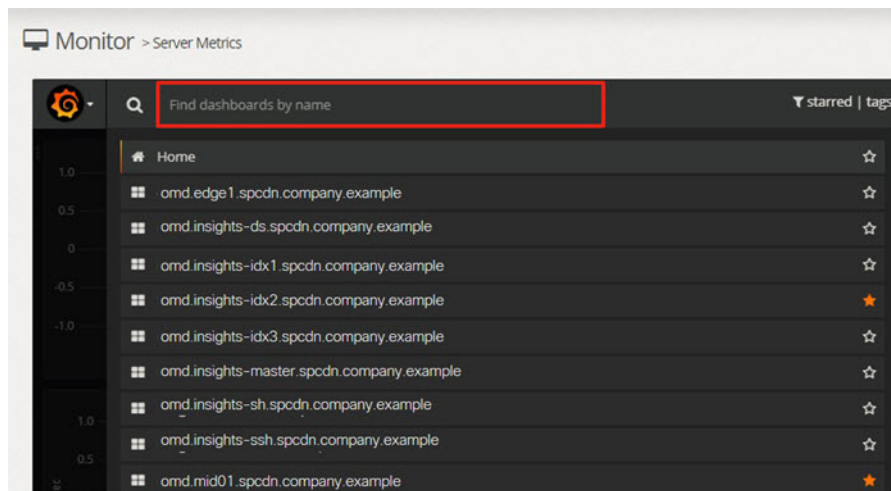
Note

The first time that you choose Monitor > Server Metrics, you will be prompted to log in to Grafana. Grafana is an open source graphical interface used for visualizing time series data through dashboards, rich graphing and alerting, and is the user interface for OMD Monitor. The default username and password will depend on the installation. For more information, please refer to the *Cisco Media Streamer and Cisco Media Broadcaster Installation and Upgrade Guide*. This login is only required once per browser. After your initial login to Grafana, you can save your username and password in the browser, which eliminates the need to authenticate each additional time you choose Monitor > Server Metrics.

The Home Dashboard that appears shows you any server dashboards that you have marked as favorite (starred) and any recently viewed server dashboards.

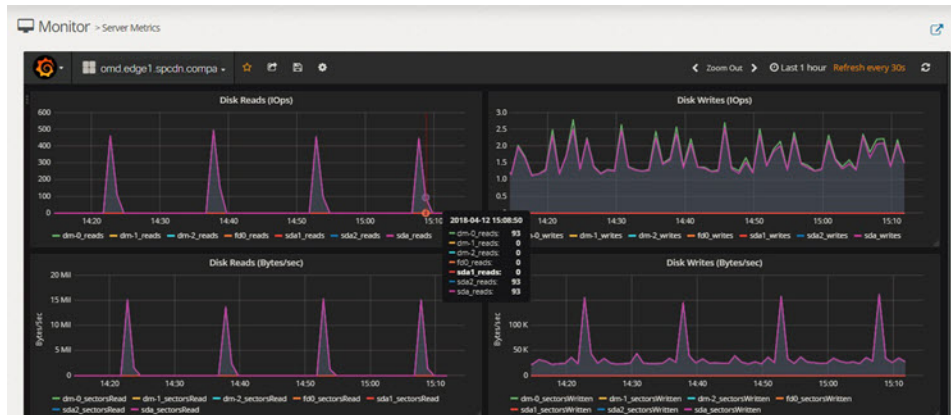
To see the system-level statistics for a specific server in the CDN environment, do one of the following:

- Click the link for the server dashboard under Starred Dashboards if it exists.
- Click the link for the server dashboard under Recently Viewed Dashboards if it exists.
- At the top of the Home dashboard, choose the server from the Home drop-down list. You can filter the list of servers by entering part of the server name in the Find Dashboard by Name text box at the top of the Home drop-down list.



From the server dashboard you can see the following information about that server:

- Disk reads based on IO per second and bytes per second
- Disk writes based on IO per second and bytes per second
- Ingress and egress bandwidth
- CPU usage and CPU load average
- Memory usage
- Disk usage

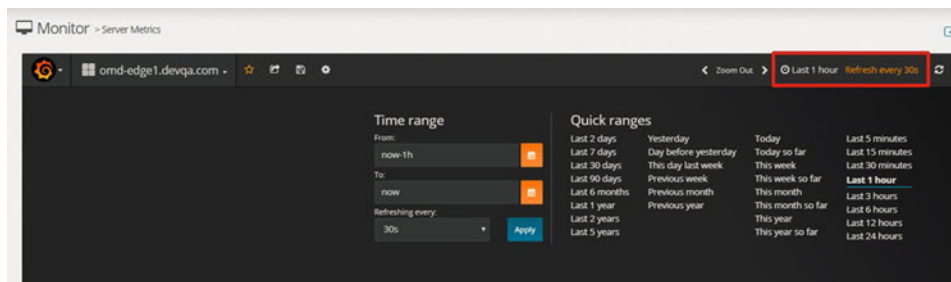


Change Time Range and Interval

By default, the last hour of statistics is displayed and are refreshed every five minutes. To see the statistics for a different time frame and optionally change the refresh rate, perform the following steps:

Procedure

- Step 1** Click the time frame icon in the tool bar. This will show the Time Range section of the dashboard at the top of the dashboard.



- Step 2** From the Time Range panel you can choose a preconfigured time range from the Quick Ranges section or you can create a custom time range in the Time Range section. To configure a custom time range, perform the following steps:
- Use the calendar icon to choose a From and To date.

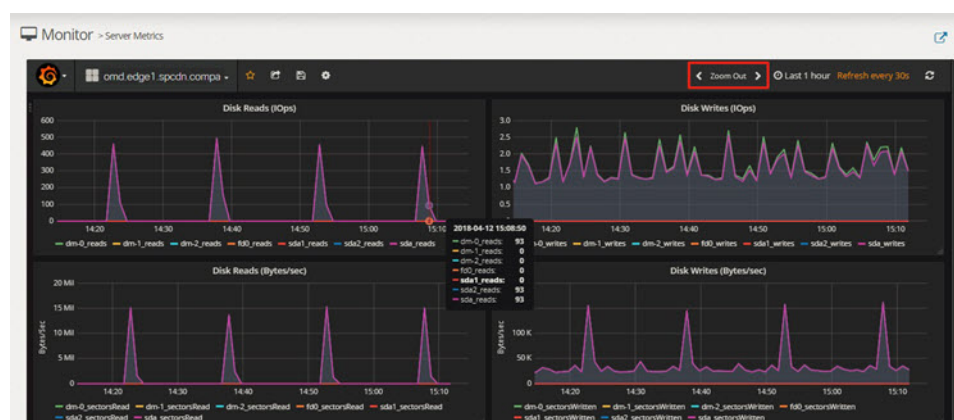
Note Optionally you can manually enter a date. If you manually enter a date, it must be in the following format: *year-month-date hours:minutes:seconds*, where *year*, *month*, and *date* are numbers. The year can be either a 2 digit or 4 digit number and the time argument is optional.
 - Optionally choose a different refresh interval from the Refreshing Every drop-down list. The default is 5 minutes.
 - Click **Apply** to apply the changes and close the Time Range panel.
- Step 3** To change the refresh interval of a preconfigured time range, perform the following steps:
- Click the time frame icon in the tool bar to expand the Time Range panel.

- From the Time Range panel, click the preconfigured time range in the **Quick Ranges** area that you want to change. This will change the time range and close the Time Range panel.
- Click the time frame icon in the tool bar again to expand the Time Range panel again.
- From the **Refreshing Every** drop-down list choose a new refresh interval and click **Apply**.

Note The new time range and refresh interval are reset to the defaults as soon as you move to a different dashboard.

Shift Time and Zoom Time

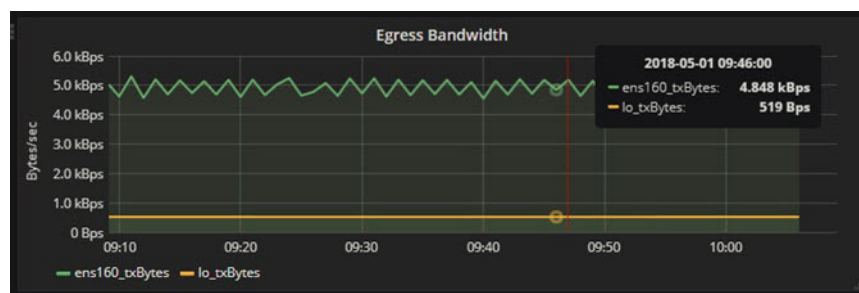
You can use the shift time (left and right arrows) and Zoom Out icons at the top of the toolbar to change the time range represented in the graphs without going into the Time Range panel.



When you click the Zoom Out icon, the time range that is displayed in the graph will double, based on the current time range that is selected. Every time you click the Zoom Out icon, it doubles the time range. For example, if the time range is set to the last 15 minutes, the first time you click the Zoom Out icon, the time range will increase to the last 30 minutes. The second time that you click the Zoom Out icon it will increase to 60 minutes.

Mouse Over

If you hover the mouse over the graph at a specific point, it will show you the value represented in the graph for that specific date and time.



Alarms

The Alarms page displays the OMD system alarms that have been raised in the OMD environment, based on configured thresholds. From the Alarms pane you can do the following:

- View the active alarms
- View the alarm history
- Configure thresholds for the alarms
- Mute alarms



Note If OMD Director is running in an HA configuration, you can view and filter all of the information available in the Alarms pages from both the Primary and Backup OMD Director instance. However, you cannot acknowledge or clear alarms, configure thresholds for alarms, or mute alarms if you are logged into an OMD Director instance running in Backup mode. If your OMD Director instance is in Detached mode because of a Director instance failure, you can view and filter all of the information available in the Alarms pages and you can also acknowledge, clear, and save alarms.

View Active Alarms

To view the OMD system alarms, choose **Monitor > Alarms**. The default tab that is displayed is the Active Alarms tab. From this tab you can see the active alarms. The Active Alarms page displays information about unresolved notifications in the OMD system. The following is an example:

Alarm Name	Server Name	Category	State	Severity	Summary	Time Raised	Mute/Unmute
tripwire_violations	omd-monitor.companyx.com	System	Cleared	High	The filesystem integrity check has detected modifications to system files	05 Jun 2018 05:46:32	Mute/Unmute
dns_resolved_status	omd-monitor.companyx.com	System	Cleared	High	The DNS resolution is failing	05 Jun 2018 05:46:37	Mute/Unmute
influx_connectivity	influx_node(s)	System	Cleared	High	None of the OMD Monitor influx nodes are reachable at this time	07 Jun 2018 05:48:15	Mute/Unmute

Each alarm will have one of the following states:

- **New:** This is an alarm that has not yet been cleared or acknowledged.
- **Cleared:** This is an alarm that has been cleared in the backend but not yet acknowledged by a user in OMD Director.
- **Acknowledged:** This is an alarm that has been acknowledged by a user in OMD Director but has not yet been cleared by the backend.



Note Alarms that are cleared in the backend and have been acknowledged by the user in OMD Director are considered resolved. Resolved alarms do not appear in the Active Alarms tab they appear in the Alarms History Tab.

Each alarm will have one of the following severities:

- **Critical:** This is an alarm that has not yet been cleared or acknowledged.
- **Warning:** This is an alarm that has been cleared in the backend but not yet acknowledged by a user in OMD Director.
- **Acknowledged:** This is an alarm that has been acknowledged by a user in OMD Director but has not yet been cleared by the backend.
- **Resolved:** An alarm is considered resolved when it is cleared in the backend and it is acknowledged by the user in OMD Director. Resolved alarms do not appear in the Active Notifications section they appear in the Alarms History section.

The Notifications column also shows whether the alarm is muted. To mute or unmute an alarms, see [Mute Alarms, on page 213](#).

Acknowledge an Alarm

To resolve a cleared alarm, which will remove it from the **Active Alarms** tab and save it in the alarms history, you must acknowledge the alarm. Follow these steps to acknowledge an alarm:



Note If you are running OMD Director in an HA configuration, you cannot acknowledge alarms from an OMD Director instance in Backup mode. You can acknowledge an alarm from an OMD Director instance in Primary or Detached mode.

Procedure

- Step 1** Click the link in the Alarm Name column for the alarm that you want to acknowledge. The Alarm Details window appears.
- Step 2** In the Alarms Details window to see a history of the Alarm before you acknowledge it, click **Show History**. If the alarm has been cleared by the backend, it will show you the date and time that the alarm was cleared.
- Step 3** Click **Yes** for Acknowledged.
- Step 4** Optionally enter a comment in the **Comment** field.
- Step 5** Click **Save**.

Alarm Details

Category: System

Severity: Critical

Entity: influx_node(s)

Current Value: 2

Description: None of the OMD Monitor Influx nodes are reachable at this time

Raised On: 07 Jun 2018 05:48:15

Acknowledged: ☒ Yes ☐ No

Comments: Enter your comments

[Show History](#) [Save](#)

Step 6 To acknowledge multiple alarms at the same time, check the check box for each alarm that you want to acknowledge and then click the Acknowledge button. To acknowledge all of the alarms, check the check box in the column header and then click the **Acknowledge** button.

Active Alarms

Alarm Name: All Category: All Severity: All Start Date: 14/02/2 12:00 A End Date: 14/06/2 11:59 P Filter Clear Filter

<input type="checkbox"/>	Alarm Name	Server Name	Category	State	Severity	Summary	Time Raised	Mute/Unmute
<input checked="" type="checkbox"/>	influx_connectivity	influx_node(s)	System	New	Critical	None of the OMD Monitor Influx nodes are reachable at this time	07 Jun 2018 05:48:15	
<input checked="" type="checkbox"/>	dns_resolved_status	omd-monitor.omdauto.com	System	New	Critical	The DNS resolution is failing	05 Jun 2018 05:46:37	
<input type="checkbox"/>	tripwire_violations	omd-monitor.omdauto.com	System	New	Critical	The filesystem integrity check has detected modifications to system files	05 Jun 2018 05:46:32	

Step 7 In the Alarm Acknowledge Comments window that appears, enter a comment and click **Save**.

Note If you acknowledge an alarm that has not been cleared by the backend, it will remain in the Active Notifications section and will show a state of Acknowledged.

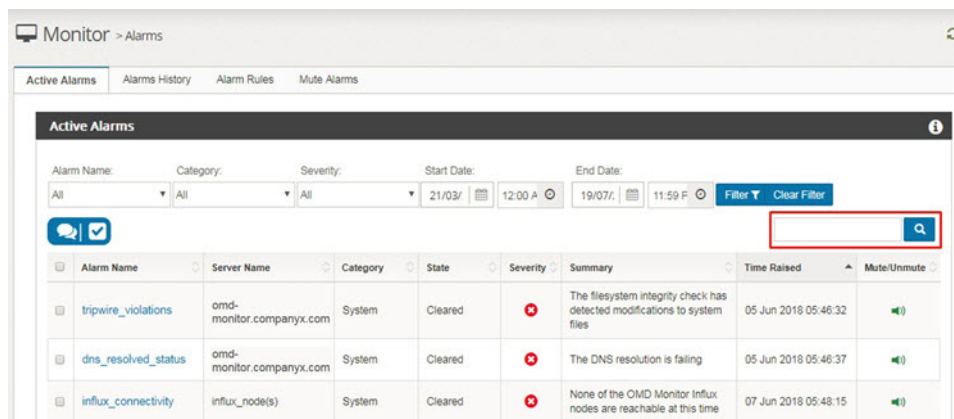
Search for an Alarm

You can also search for alarms by entering text in the search field at the top of the Active Notifications section. The text that you enter is searched for in all of the columns. To search based on the severity of the alarm, enter the severity level you are looking for: critical, warning, informational, or debug.



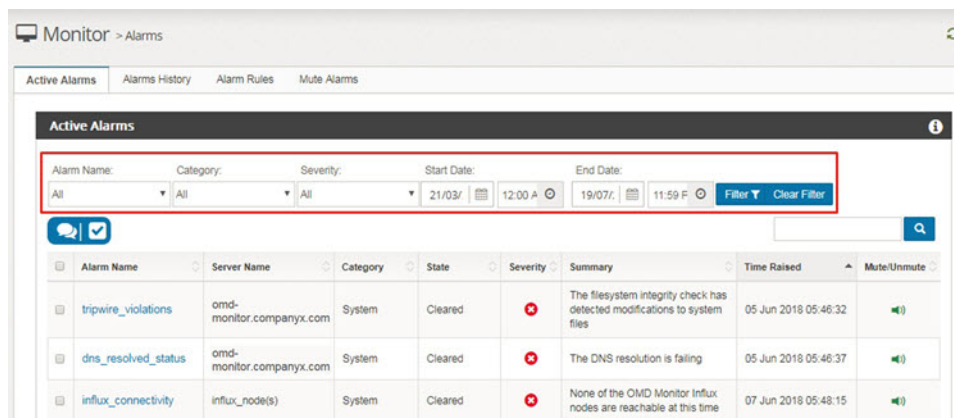
Note Searches are case-sensitive.

Filter an Alarm



Filter an Alarm

You can also filter the list of alarms by name, category, severity, and date. Enter values in the filter toolbar at the top of the **Active Alarms** tab and click **Filter**.



Alarms History

The **Alarms History** tab from the Monitor > Alarms page displays resolved alarms. An alarm is considered resolved when it is cleared by the backend and acknowledged by a user in OMD Director. The following is an example:

Alarms History							
Alarm Name:	All	Category:	All	Severity:	All	Start Date:	07/02/2018
						End Date:	09/03/2018
						Filter	
						Show	10 entries
Alarm Name	Category	Severity	Raised On	Details	Last Updated By	Last Updated On	
bond_interface	system	warning	07 Feb 2018 03:08:25	Bond interface is not configured	omdadmin	07 Feb 2018 22:43:53	
bond_interface	system	warning	07 Feb 2018 03:08:25	Bond interface is not configured	omdadmin	07 Feb 2018 22:43:53	
bond_interface	system	warning	07 Feb 2018 03:08:26	Bond interface is not configured	omdadmin	07 Feb 2018 22:43:53	
bond_interface	system	warning	07 Feb 2018 03:08:38	Bond interface is not configured	omdadmin	07 Feb 2018 22:43:53	
bond_interface	system	warning	07 Feb 2018 03:09:57	Bond interface is not configured	omdadmin	07 Feb 2018 22:43:53	
bond_interface	system	warning	07 Feb 2018 03:09:13	Bond interface is not configured	omdadmin	07 Feb 2018 22:43:53	
bond_interface	system	warning	07 Feb 2018 03:09:16	Bond interface is not configured	omdadmin	07 Feb 2018 22:43:53	
bond_interface	system	warning	07 Feb 2018 03:09:17	Bond interface is not configured	omdadmin	07 Feb 2018 22:43:53	
bond_interface	system	warning	07 Feb 2018 22:43:23	Bond interface is not configured	sathya	07 Feb 2018 23:38:28	

Filter Alarms History

By default, the **Alarms History** tab shows the resolved alarms for the past 60 days, but you can change that date range. You can also filter the list of alarms by name, category, and severity. To change the date range that is displayed or to filter based on the alarm name, category, or severity, enter the values in the toolbar at the top of the **Alarms History** tab and click **Filter**.

Alarms History							
Alarm Name:	All	Category:	All	Severity:	All	Start Date:	25/02/2018
						End Date:	27/03/2018
						Filter	
						Show	10 entries
Alarm Name	Category	Severity	Raised On	Details	Last Updated By	Last Updated On	
keep_alive	system	Warning	26 Feb 2018 04:34:09	The system is not reachable by the OMD Monitoring Server and is not responding to keepalive events for > 90s	omdadmin	22 Mar 2018 05:24:57	
keep_alive	system	Warning	26 Feb 2018 04:35:39	The system is not reachable by the OMD Monitoring Server and is not responding to keepalive events for > 90s	omdadmin	22 Mar 2018 05:24:57	

Configure Alarm Rules

All notifications are generated based on the rules that are configured for the alarms. These rules are configured by editing or creating a new alarm rule. An alarm can have multiple rules configured for it to send notifications about different severity levels. To view the existing alarm rules and to create new alarm rules, from the Monitor > Alarms page, click the **Alarm Rules** tab.



Note

The [OMD Director Alarms and Remediation, on page 289](#) describes the default checks and thresholds that are run by the monitor-client service on the OMD Monitor clients that you can view and deactivate in OMD Director from the Alarms Rules tab.

**Note**

If you are running OMD Director in an HA configuration, you can only make configuration changes from the Primary OMD Director instance. To perform the steps in this chapter, make sure you are logging into the Primary OMD Director instance.

Monitor > Alarms

Active Alarms Alarms History **Alarm Rules** Mute Alarms

Alarm Rules

Create New

Search Thresholds

Name	Category	Entity Family	Threshold	Severity	Description	Action
alarms_db_backup	system	director-worker	2	critical	Alarms db backup job has been failed	Deactivate
bond_interface	system	cache_nodes	2	critical	Bond interface is currently down	Deactivate
cpu_usage	system	common	50	warning	The percentage of CPU utilized is above configured threshold level	Deactivate
cpu_usage	system	common	80	critical	The percentage of CPU utilized is above configured threshold level	Deactivate
disk_drives_count	system	cache_nodes	2	critical	The number of disk drive partitions currently available is less than the original number of partitions	Deactivate
disk_usage	system	common	1	warning	The percentage of disk utilized is above configured threshold level and the file system is nearing its full capacity	Deactivate
disk_usage	system	common	95	critical	The percentage of disk utilized is above configured threshold level and the file system is nearing its full capacity	Deactivate

Create New Alarm Rule

Follow these steps to create a new alarm rule for an alarm:

Procedure

Step 1

From the Monitor > Alarms page, click **Create New**. The Add New Alarm Rule window appears.

- Step 2** In the **Alarm Name** field, choose an alarm from the auto-populated alarms list for which you want to create the rule.
- Step 3** From the **Category** drop-down list, choose the alarm category. Currently, the only category available is System.
- Step 4** From the **Entity Family** drop-down list, choose the entity family to which this rule will apply. This determines to which CDN elements this rule will apply.
- Step 5** From the **Alert Condition** drop-down list choose one of the following:
- Less than
 - Greater than
 - Equal to
 - Less than or Equal to
 - Greater than or Equal to
- Step 6** In the **Threshold** field, enter the value that should trigger the notification based on the selected Alert Condition.
- Step 7** In the **Description** field, enter a description for the rule. It is helpful if the description includes information about the threshold value that needs to be reached to trigger the notification.
- Step 8** From the **Severity** drop-down list, choose the severity level to assign to this rule. The levels are: Critical, Warning, Informational, and Debug.
- Step 9** In the **Remediation Steps** field, optionally enter a URL in this field that points to instructions for the remediation of the alarm. This URL will be included in any email or SMS notifications that are triggered by this rule. For more information on how to configure which alarms you would like to receive notifications for and how to receive them, see [Notification Settings, on page 253](#).
- Step 10** Click **Add** to save the new alarm rule. The following is an example:

Add New Alarm Rule

Alarm Name *

Category *

Entity Family *

Alert Condition

Threshold *

Description *

Severity

Remediation Steps

Edit an Alarm Rule

To edit an existing alarm rule, click the name of the alarm rule you want to edit.

Monitor > Alarms

Active Alarms Alarms History Alarm Rules Mute Alarms

Alarm Rules

[Create New](#)

Name	Category	Entity Family	Threshold	Severity	Description	Action
alarms_db_backup	system	director-worker	2	critical	Alarms db backup job has been failed	Deactivate
bond_interface	system	cache_nodes	2	critical	Bond interface is currently down	Deactivate
cpu_usage	system	common	50	warning	The percentage of CPU utilized is above configured threshold level	Deactivate
cpu_usage	system	common	80	critical	The percentage of CPU utilized is above configured threshold level	Deactivate
disk_drives_count	system	cache_nodes	2	critical	The number of disk drive partitions currently available is less than the original number of partitions	Deactivate
disk_usage	system	common	1	warning	The percentage of disk utilized is above configured threshold level and the file system is nearing its full capacity	Deactivate
disk_usage	system	common	95	critical	The percentage of disk utilized is above configured threshold level and the file system is nearing its full capacity	Deactivate

The edit panel for the rule expands. From this window you can edit the description, alert condition and threshold value, severity, and link for remediation steps. When you are done making changes, click the **check mark** icon. To cancel the changes, click the **x** icon.

[cpu_usage](#) system common 50 warning The percentage of CPU utilized is above configured threshold level [Deactivate](#)

cpu_usage

Description *

The percentage of CPU utilized is above configured threshold level

Alert Condition *

Greater than Or Equal to 80

Severity *

critical

Remediation Steps

Remediation Steps Link

Deactivate an Alarm Rule

To deactivate an alarm rule, which will prevent notifications from being triggered for this rule until it is reactivated, click the **Deactivate** link, which appears in the last column of the rule.

Monitor > Alarms

Active Alarms Alarms History Alarm Rules Mute Alarms

Alarm Rules

Create New

Search Thresholds

Name	Category	Entity Family	Threshold	Severity	Description	Action
alarms_db_backup	system	director-worker	2	critical	Alarms db backup job has been failed	Deactivate
bond_interface	system	cache_nodes	2	critical	Bond interface is currently down	Deactivate
cpu_usage	system	common	50	warning	The percentage of CPU utilized is above configured threshold level	Deactivate
cpu_usage	system	common	80	critical	The percentage of CPU utilized is above configured threshold level	Deactivate
disk_drives_count	system	cache_nodes	2	critical	The number of disk drive partitions currently available is less than the original number of partitions	Deactivate
disk_usage	system	common	1	warning	The percentage of disk utilized is above configured threshold level and the file system is nearing its full capacity	Deactivate
disk_usage	system	common	95	critical	The percentage of disk utilized is above configured threshold level and the file system is nearing its full capacity	Deactivate

Mute Alarms

From the **Mute Alarms** tab you can suspend notifications for specific alarms on specific servers for a configured period of time. You can also see which alarms for which servers are currently muted.

Active Alarms Alarms History Thresholds Mute Alarms

Mute Alarms

Alarm Names * All

-- Select Alarms --

Server Names * All

-- Select Servers --

Comments *

Comments

Mute Till*

Select a date

Select time

Mute

Currently Muted Alarms

Alarm Name	Server List	Muted By	Comments	Mute Till
cpu_usage	edge-cache2.company.example	omdadmin	test	09 Mar 2018 05:40 PM

5 Minutes Snooze Un-Mute

To mute an alarm, perform the following steps:

Procedure

Step 1

From the **Alarm Names** drop-down list, choose the alarm that you want to mute. You can choose multiple alarms if needed. To select additional alarms, click in the **Alarm Names** field and the **Alarm Names** drop-down list will appear again allowing you to select an additional alarm.

- Step 2** From the **Server Names** drop-down list, choose the server for which you want to mute the alarm. You can choose multiple servers if needed. To select additional servers, click in the **Server Name** field, and the Server Names drop-down list will appear again allowing you to select an additional server.
- Step 3** Enter a description.
- Step 4** In the Mute Till field, use the calendar and time icons to choose when notifications for the alarm should resume.
- Step 5** When you are finished making changes, click **Mute** to mute the alarms.
- Step 6** The muted alarms will now appear in the **Currently Muted Alarms** area.

The screenshot shows the 'Mute Alarms' interface. At the top, there are tabs for 'Active Alarms', 'Alarms History', 'Thresholds', and 'Mute Alarms'. The 'Mute Alarms' tab is selected. Below the tabs, there is a form titled 'Mute Alarms'. The form has two rows of input fields. The first row has 'Alarm Names' with a dropdown menu (currently showing 'All') and 'Server Names' with a dropdown menu (currently showing 'All'). Below these are two more dropdown menus: '-- Select Alarms --' and '-- Select Servers --'. The second row has a 'Comments' field and a 'Mute Till' field. The 'Mute Till' field has a calendar icon and a clock icon. To the right of the 'Mute Till' field is a 'Mute' button. Below the form is a section titled 'Currently Muted Alarms' which contains a table of muted alarms. The table has columns for 'Alarm Name', 'Server List', 'Muted By', 'Comments', and 'Mute Till'. There is one row in the table with the following data: 'cpu_usage', 'edge-cache2.company.example', 'omdadmin', 'test', and '09 Mar 2018 05:40 PM'. To the right of the table is a '5 Minutes' dropdown menu and two buttons: 'Snooze' and 'Un-Mute'.

Alarm Name	Server List	Muted By	Comments	Mute Till
cpu_usage	edge-cache2.company.example	omdadmin	test	09 Mar 2018 05:40 PM

From the **Currently Muted Alarms** area you can automatically unmute the alarm without waiting for the Mute Till date to be reached by clicking **Un-Mute**. You can also extend the **Mute Till** date by clicking **Snooze**. This will extend the Mute Till date by the value selected in the time drop-down list that appears before the Snooze button.



CHAPTER 13

OMD Insights

OMD Insights provides deep analytics of the OMD environment that is based on mining data from the HTTP log entries of each Traffic Server. This enables OMD Insights to provide comprehensive CDN operational analytics, including analysis of throughput, utilization, and efficiency of video content delivery. It also provides insight into viewer trends. OMD Insights captures real-time data from CDN logs to provide simplified access to actionable data and analytics, so that operation teams can proactively monitor activity and take action before problems occur. It also delivers comprehensive dashboards and trending reports that capture valuable information for CDN capacity planning and delivery optimization. In general OMD Insights provides transaction logging and application level analytics.

This chapter provides an overview of the information that is available from the Insights page in OMD Director.

This chapter includes the following topics:

- [Insights Page Overview](#), on page 215
- [Insights Overview Tab](#), on page 216
- [Insights Trends Tab](#), on page 222
- [Insights Reports Tab](#), on page 223
- [Insights Analytics Tab](#), on page 224
- [Insights Monitor Tab](#), on page 226
- [Insights Alerts Tab](#), on page 227
- [Insights Custom Dashboards Tab](#), on page 229
- [Insights Search Tab](#), on page 233

Insights Page Overview

You can access OMD Insights directly from OMD Director using the Insights menu in the navigational panel. All of the information provided in the different dashboards, graphs, and charts is pulled from the different servers that are part of the OMD Insights deployment. Each tab contains predefined charts, graphs, and reports to analyze key application level analytics of the CDN deployment.

The Insights page contains the following tabs:

- **Overview:** Contains key charts to help provide a quick overview of the overall performance of the CDN network.
- **Trends:** It provides long time-range and specific day analysis of individual CDN metrics.
- **Reports:** Combines the KPIs and metrics for a set of predefined time-ranges into a tabular format.

- **Analytics:** Contains a set of tools for granular analysis.
- **Monitor:** Contains charts that present CDN metrics in real-time.
- **Alerts:** From this tab you can create and view thresholds that trigger email alerts.
- **Custom Dashboard:** From this tab you can view, create, and delete custom dashboards.
- **Search:** From this tab you can perform custom searches of the analytics transaction database.

From the main Insights page you can also create thresholds to trigger alerts, create custom dashboards, and create custom searches of the analytics database.



Note To change the formatting of numeric data to use a period to separate groups of thousands (instead of a comma), and to use commas for decimal places (instead of periods), enter “| converttoitalianformat” (without the quotes) at the end of any search query.



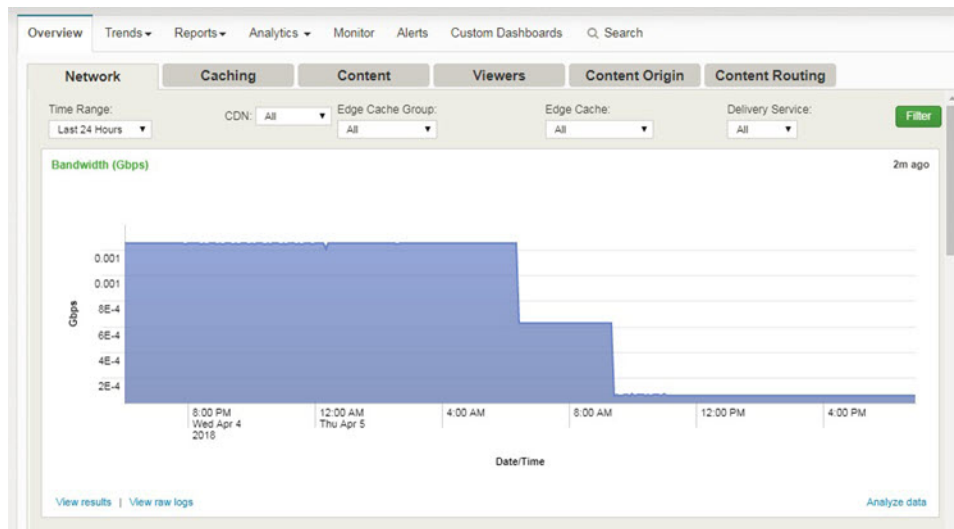
Note OMD Insights is based on Splunk, which uses buckets to store the indexed data that is used to create the charts and tables provided by OMD Insights. If you notice a gap in data in any of the charts or tables, or if you notice a sudden spike in any of the values, this indicates that a summary bucket may be missing or duplicated. To resolve this issue, you need to delete the existing summary index buckets for the appropriate time frame and recreate them. For information on how to delete the existing summary index buckets and recreate them, see [Recreate OMD Insights Summary Index Buckets, on page 275](#).

Insights Overview Tab

When you choose Insights from OMD Director, a page is launched to the OMD Insights product and the Overview tab is displayed. The Overview tab contains key charts to help provide a quick overview of the overall performance of the CDN network. Metrics for these charts are derived from summarized data: the incoming logs are indexed and searches are scheduled to create summarized data from these logs. Content for these charts is limited to a maximum of seven days.

The Overview tab contains the following sub tabs:

- Network
- Caching
- Content
- Viewers
- Content Origin
- Content Routing



Filtering Content

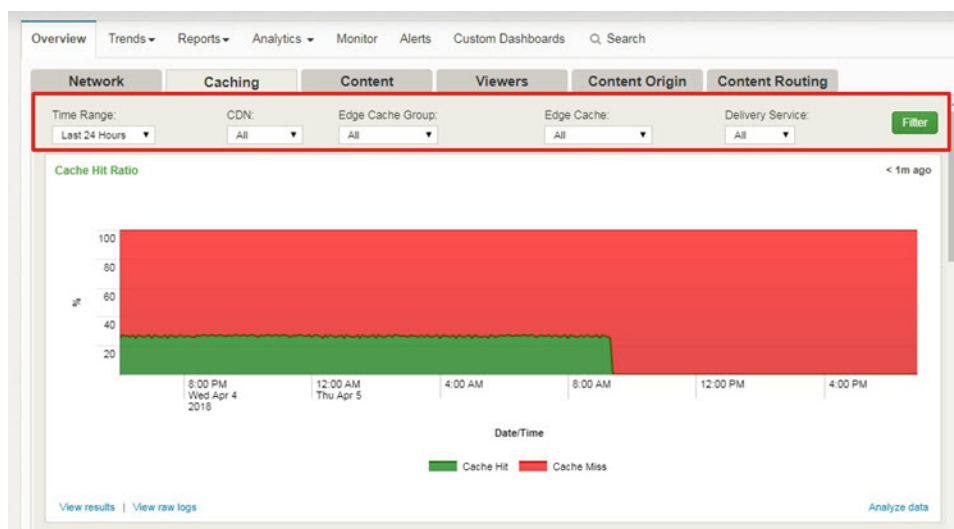
You can change the time range for the content that is displayed in the charts and tables that are available from the Overview tab by using the **Time Range** drop-down list. Which Time Range options you have will depend on the tab, however, content for these charts is limited to a maximum of seven days. To see data across a longer period of time, use the charts and reports available from the **Trends** tab.

You can also filter the content of the charts and graphs that are displayed on the Overview sub tabs based on additional criteria, such as Edge Cache Group and Delivery Service, by using the Filter toolbar at the top of the tab. What criteria you can filter on will vary depending on which tab you are on.



Note

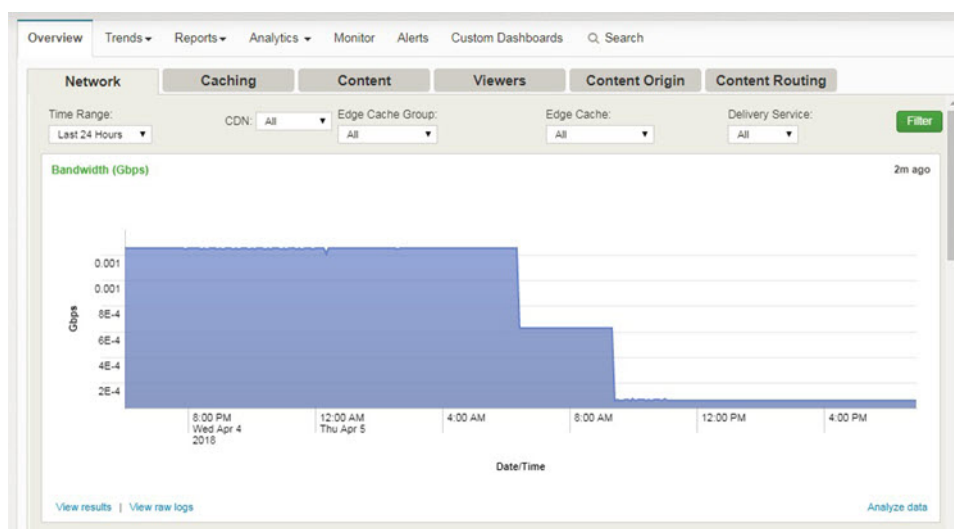
After you make any changes to any of the criteria on the Filter toolbar, you must click **Filter** for those changes to take affect.



Network Tab

The network tab contains the following charts:

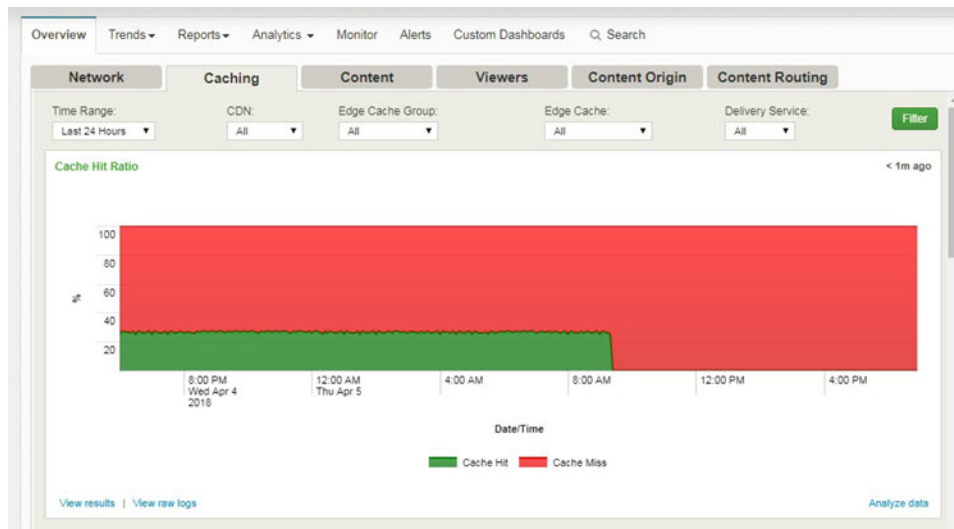
- **Bandwidth:** This chart provides the average delivered bandwidth in Gigabytes per second (Gbps) across the selected time range.
- **Bandwidth Distribution:** This chart provides the average delivered bandwidth broken into different bandwidth ranges across the selected time range.
- **Volume Delivered by Platform:** This chart provides the average delivered bandwidth broken down by the platform across the selected time range.
- **Volume Delivered by Protocol:** This chart provides the average delivered bandwidth broken down by the protocol across the selected time range.



Caching Tab

The **Caching** tab contains the following charts:

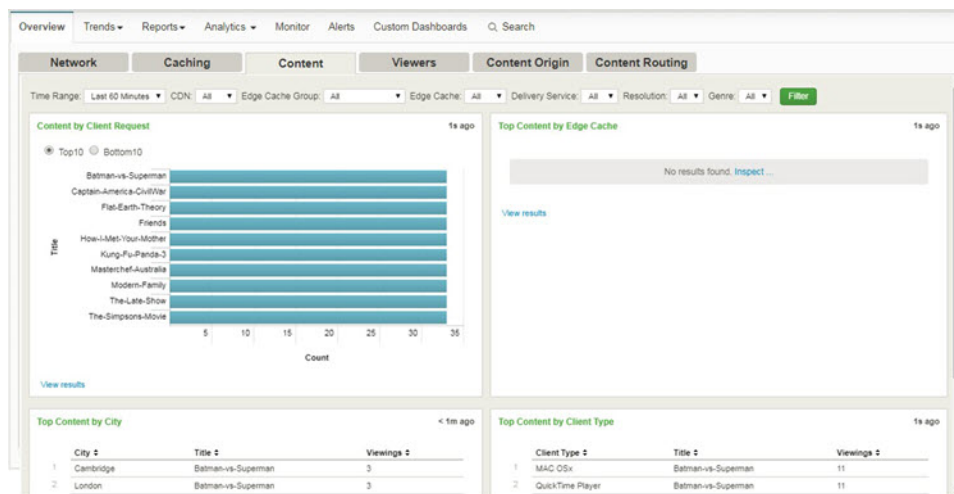
- **Cache Hit Ratio:** This chart shows the cache hit and miss ratio for all caches, across the specified time range.
- **Cache Hit Percentage by Cache Groups:** This chart shows the percentage of cache hits and misses received broken down by the cache groups, across the specified time range.
- **Cache Hit Percentage by Delivery Service:** This chart shows the percentage of cache hits and misses received broken down by the Delivery Service, across the specified time range.
- **Cache Hit Percentage by Protocol:** This chart shows the percentage of cache hits and misses received broken down by the protocols, across the specified time range.



Content Tab

The **Content** tab contains the following chart and tables:

- **Content by Client Request:** This chart shows the count of the top or bottom 10 shows viewed, based on your selection, and shows the number of viewings for each of these shows, across the selected time range.
- **Top Content by Edge Cache:** This table lists the title of the top content by edge device, including the number of viewings, across the selected time range.
- **Top Content by City:** This table lists the top content count per city, including the number of viewings, across the selected time range.
- **Top Content by Client Type:** This table lists the top content count per client type, including the number of viewings, across the selected time range.



Viewers Tab

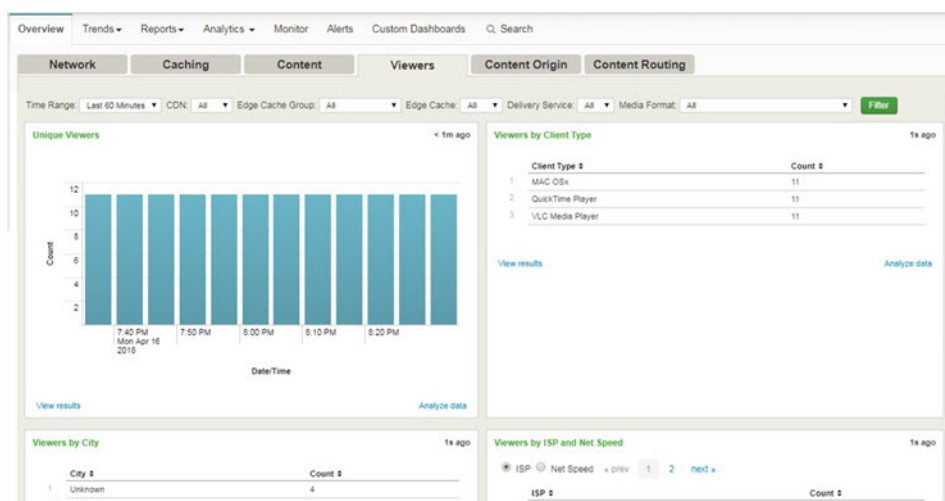
The **Viewers** tab contains the following chart and table:

- **Unique Viewers:** This chart shows the total number of unique viewer across the selected time range.
- **Viewers by Client Type:** This report lists the total number of unique viewer, broken down by client type, across the selected time range.
- **Viewers by City:** This report lists the total number of unique viewers broken down by city, across the selected time range.
- **Viewers by ISP and Net Speed:** This report lists the total number of unique viewers broken down by ISP or Net Speed, based on your selection, across the selected time range.



Note

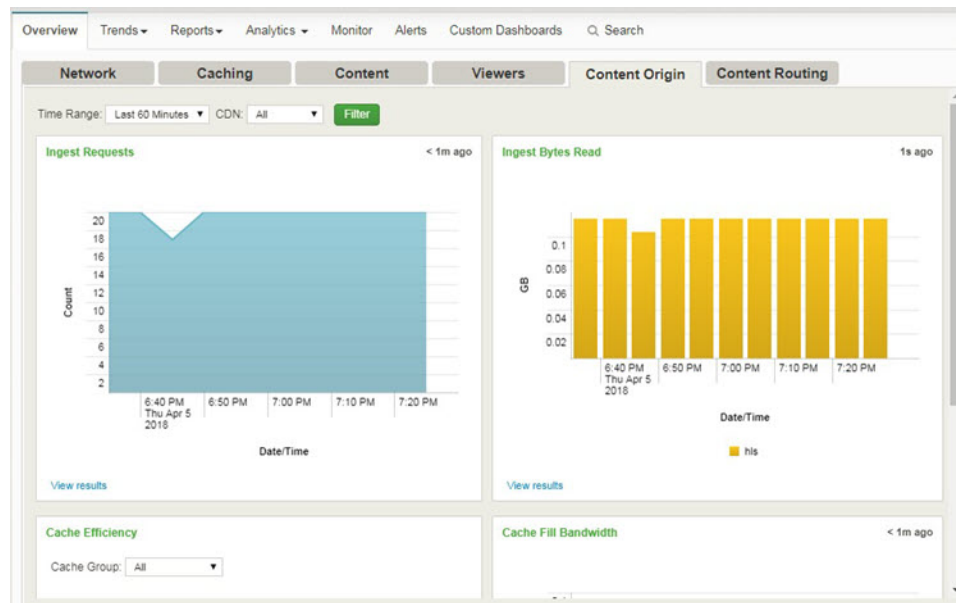
This chart only shows results if you have purchased a license for the MaxMind ISP and Connection Type databases.



Content Origin Tab

The **Content Origin** tab contains the following charts:

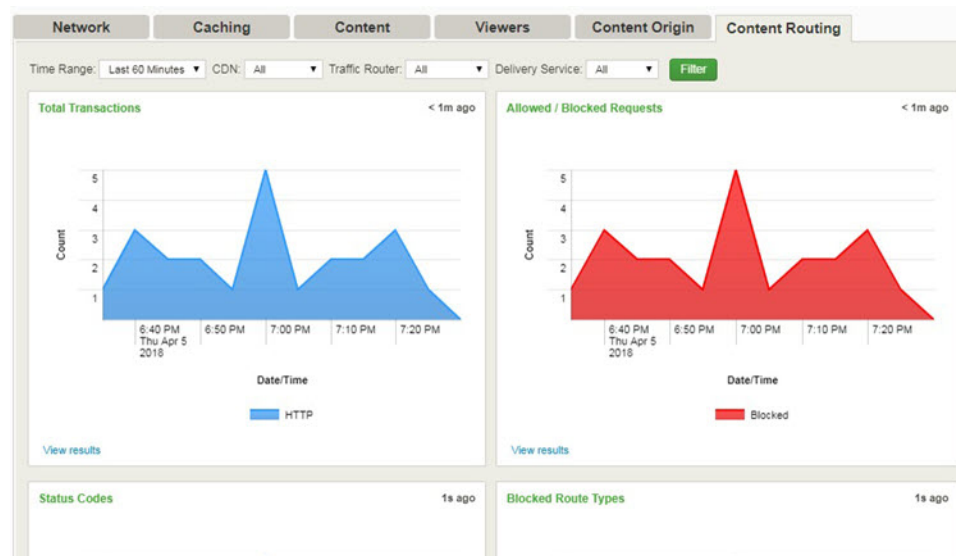
- **Ingest Requests:** This chart shows the total number of ingest requests across the selected time range.
- **Ingest Bytes Read:** This chart shows the total number of ingest bytes read across the selected time range, broken down by protocol.
- **Cache Efficiency:** This chart shows the ratio of Mid cache hits and misses for cache-fill requests from the edge caches across the selected time range.
- **Cache Fill Bandwidth:** This chart show the peak inbound bandwidth from Mid caches to Edge caches across the selected time range.



Content Routing Tab

The following charts are available from the Content Routing tab:

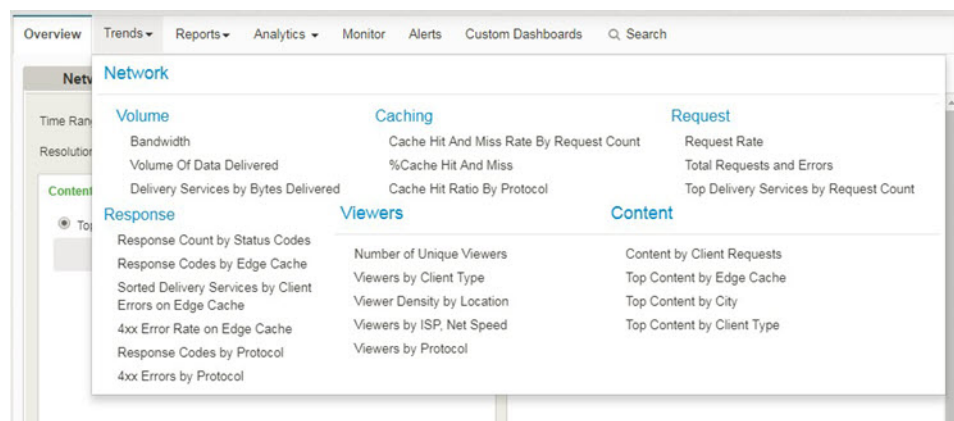
- Total Transactions
- Allowed/Blocked Requests
- Status Codes
- Blocked Route Types
- Blocked Channels



Insights Trends Tab

The Trends tab enables you to view historical data (data that is older than 7 days), in a graph and tabular format. From the Trend tab you can view information about the following items:

- Network Volume
- Network Caching
- Network Requests
- Network Responses
- Viewers
- Content



Filtering Content

You can change the time range for the content that is displayed in the charts and tables available from the Trends tab by using the Time Range drop-down list. For the charts and reports that are available from the Trends tab, you can choose from 7 days, 30 days, 90 days, 365 days, or create a custom time range.

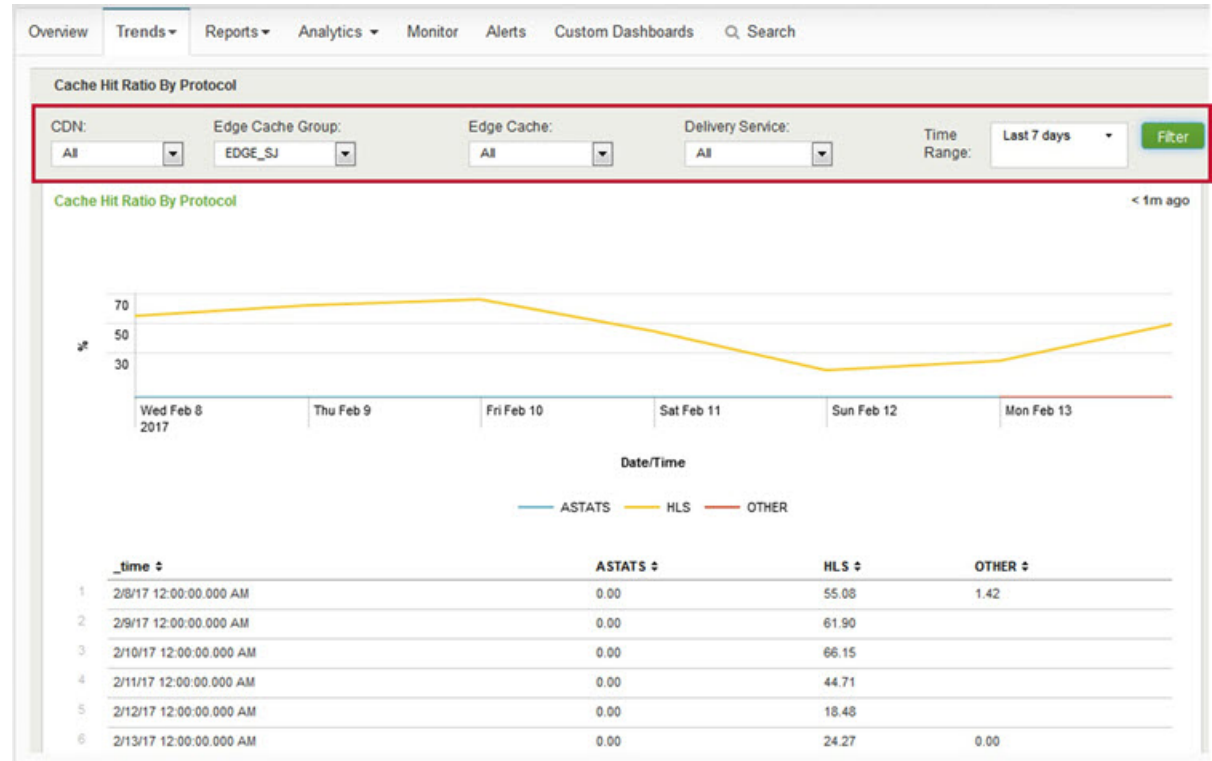
You can also filter the content of the charts and graphs that are displayed based on additional criteria, such as Edge Cache Group and Delivery Service, by using the Filter toolbar at the top of the tab. What criteria you can filter on will vary depending on which tab you are on.



Note

After you make any changes to any of the criteria on the Filter toolbar, you must click the Filter button for those changes to take affect.

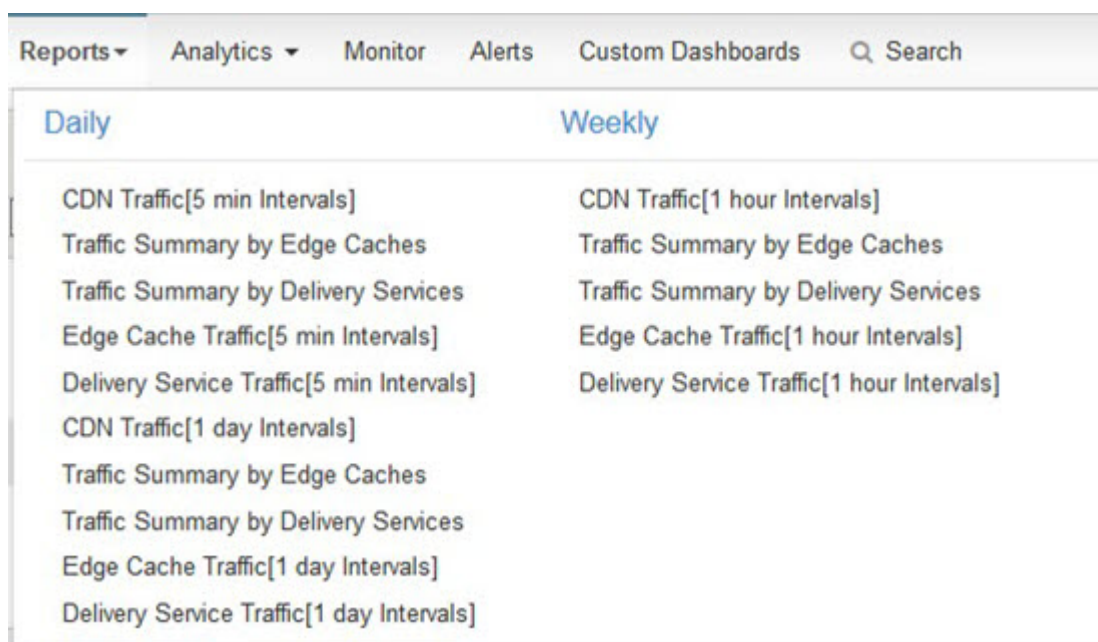
Figure 5: Filtering Trend Charts



Insights Reports Tab

The Reports tab enables you to run preconfigured Daily and Weekly reports that show information about the traffic in the CDN network. The Daily reports show information for the current day and the Weekly reports show information for the last seven days.

You cannot change the time range of these reports, but the Edge Cache Traffic reports allow you to filter on CDN and Edge Cache and the Delivery Service reports allow you to filter on CDN and Delivery Service.



The following is an example of the **Reports** tab, which displays the weekly CDN traffic:

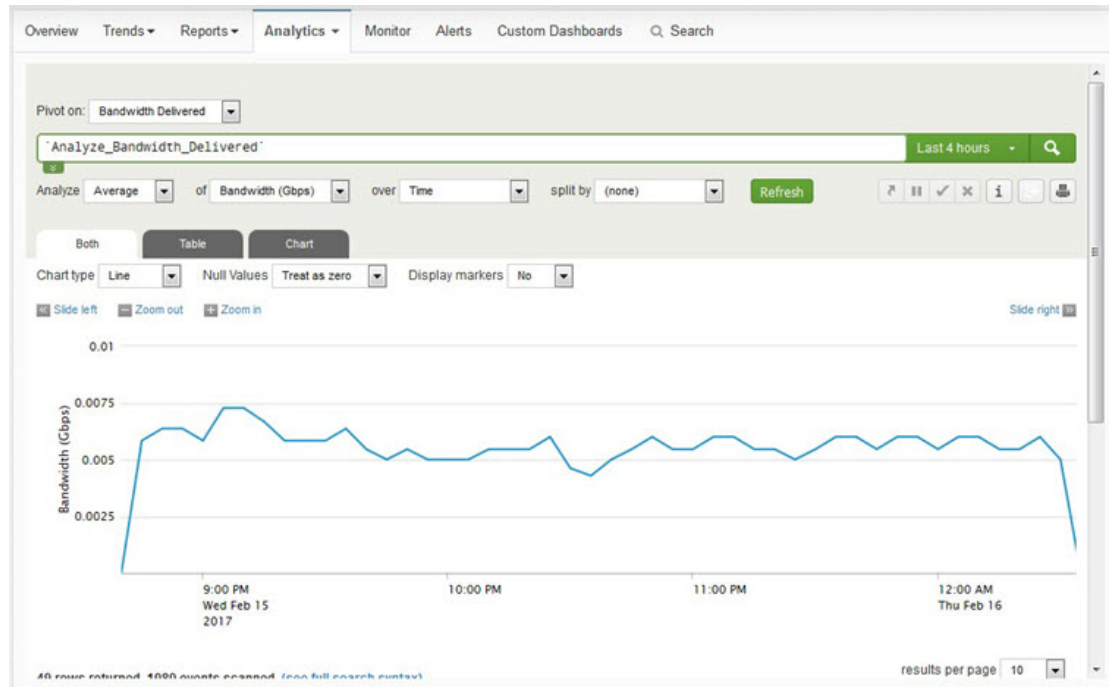
The screenshot shows the 'Reports' tab in the OMD Insights interface, displaying a table of weekly CDN traffic data. The table has the following columns: '_time', 'Bandwidth (Gbps)', 'Total GB Delivered', 'Total GB Ingested', 'Cache Hits', 'Cache Misses', 'Cache Hit Ratio', 'Total Requests', and '4xx Errors'. The data is presented in a table with 10 rows, showing traffic data for the week of 2/9/17. The table is titled 'CDN Traffic' and includes a 'View results' link at the bottom.

	_time	Bandwidth (Gbps)	Total GB Delivered	Total GB Ingested	Cache Hits	Cache Misses	Cache Hit Ratio	Total Requests	4xx Errors
1	2/9/17 12:00:00.000 AM	0.05	21.61	7.66	9173	9360	49.50	21315	2772
2	2/9/17 1:00:00.000 AM	0.05	21.51	7.56	9385	9087	50.81	21200	2719
3	2/9/17 2:00:00.000 AM	0.05	21.69	7.75	9213	9286	49.80	21238	2727
4	2/9/17 3:00:00.000 AM	0.05	21.75	7.57	9342	9246	50.26	31703	13099
5	2/9/17 4:00:00.000 AM	0.05	20.76	7.17	8272	9054	47.74	20050	2721
6	2/9/17 5:00:00.000 AM	0.05	20.63	6.97	7945	9102	46.61	24955	7908
7	2/9/17 6:00:00.000 AM	0.05	20.74	7.02	7917	9155	46.37	19792	2720
8	2/9/17 7:00:00.000 AM	0.05	20.74	7.01	7944	9136	46.51	19803	2723
9	2/9/17 8:00:00.000 AM	0.05	20.33	7.06	7839	9167	46.10	19753	2745
10	2/9/17 9:00:00.000 AM	0.05	20.88	7.69	7761	9483	45.01	19969	2721

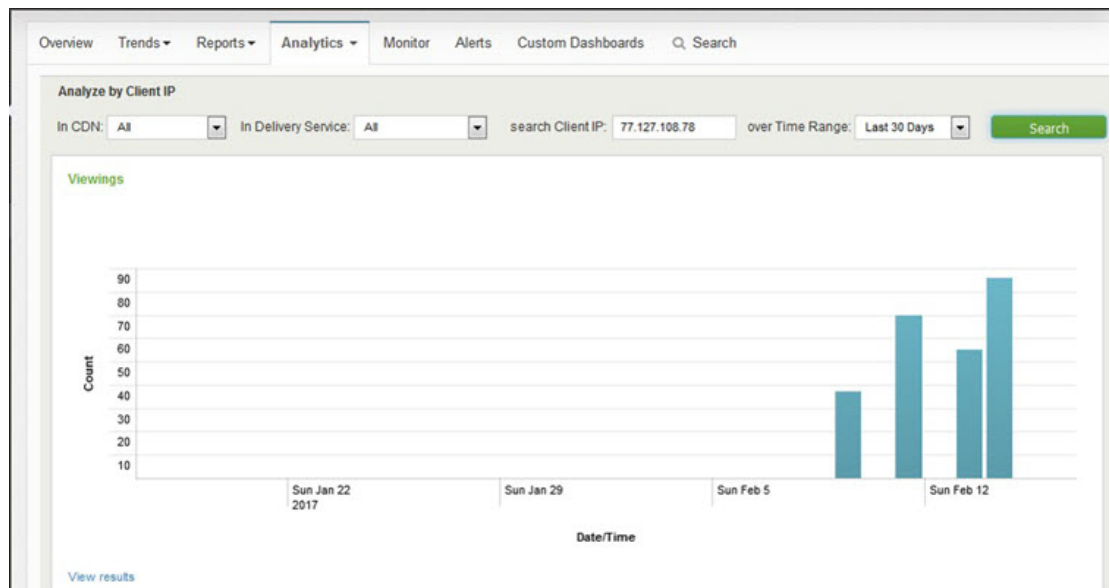
Insights Analytics Tab

The Analytics tab enables the user to analyze the CDN metrics at a granular level. The Analytics tab contains the following options:

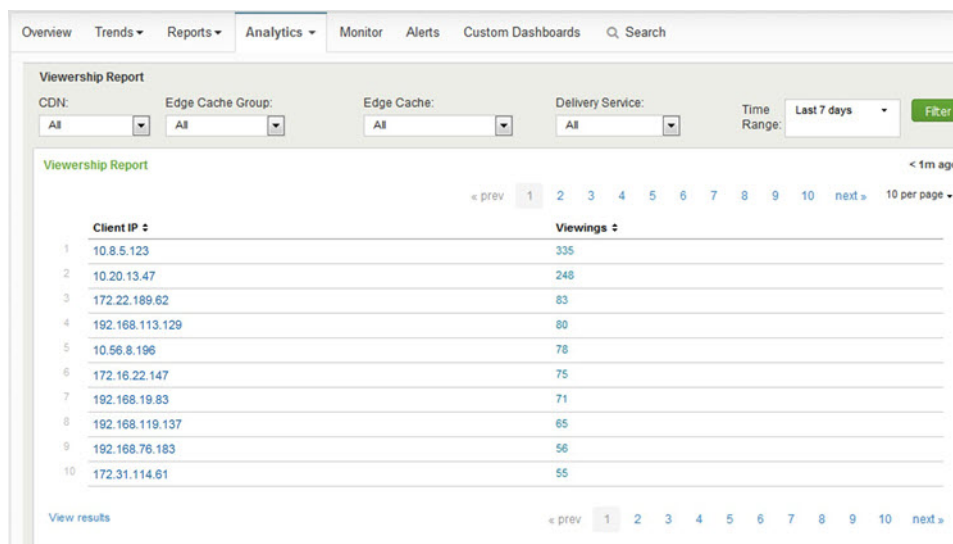
- **Analyze Metrics with Pivoting:** This option enables you to choose which data you want to pivot on and how you want to break down the data, for example based on time or edge cache. You can also filter the information to display based on the Top X or Bottom X values returned and you to change the chart type to use when displaying the chart.



- **Analyze by Client IP:** This option enables you to search for a specific client IP address to see how many viewings this client has had across the time frame selected. You must enter the exact client IP address that you want to view, you cannot enter subnets or wildcards. You can filter this report based on CDN and Delivery Service.



- **Viewership Report:** This option displays a report that shows all of the unique client IP addresses that have requested content for the selected time range and the number of viewings for that client. You can filter this report based on CDN, Edge Cache Group, Edge Cache, and Delivery Service.

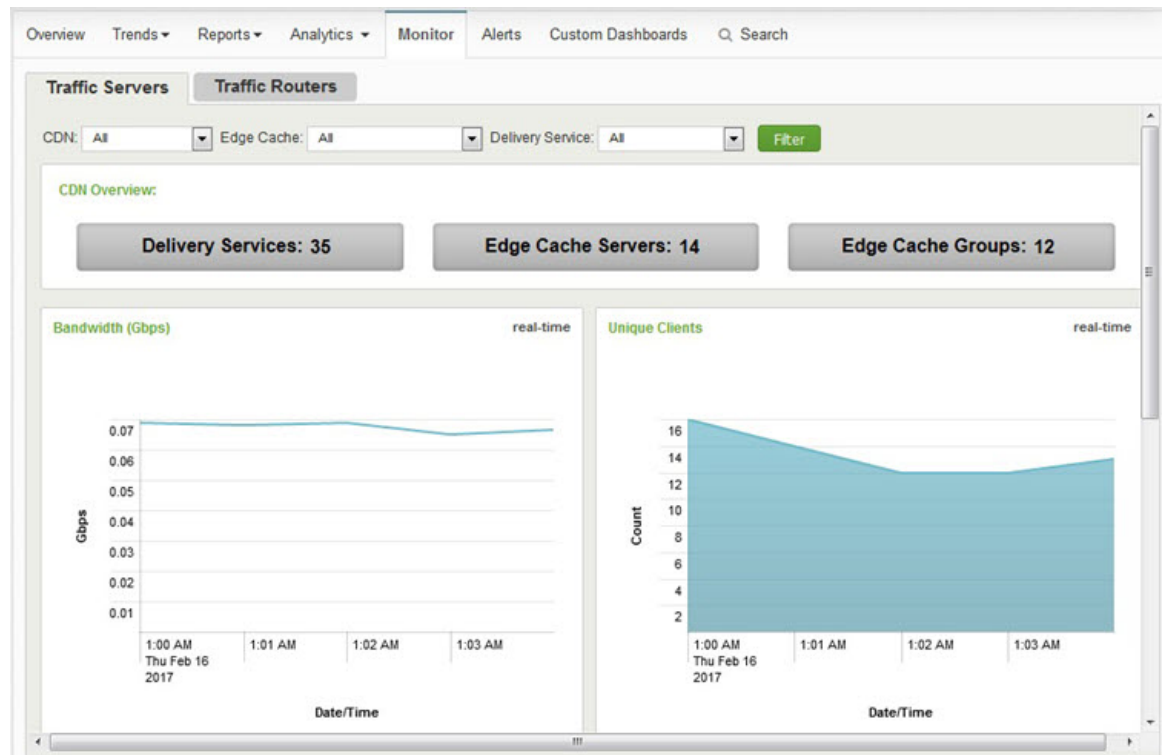


Insights Monitor Tab

The graphs that are available from the Insights Monitor tab derive their metrics using raw logs. The incoming logs are consumed as-is to plot the metrics, therefore the graphs available from the Monitor tab present CDN metrics in near real-time.

From the Traffic Servers sub tab, you can view graphs that show information about the bandwidth, clients, cache hit ratios, and responses for the Traffic Servers. This tab also shows the total number of Delivery Services, Edge cache servers, and Edge cache groups in the CDN.

From the Traffic Routers tab you can view the total transactions, blocked transactions, and responses for the Traffic Routers.



Filtering Content

You can filter the content of the charts and graphs that are displayed from the Monitor sub tabs based on additional criteria. From the Traffic Servers sub tab, you can filter the content based on CDN, Edge Cache, and Delivery Service by using the Filter toolbar at the top of the tab. From the Traffic Routers tab, you can filter the content based on CDN, Traffic Router, and Delivery Service by using the Filter toolbar at the top of the tab.



Note After you make any changes to any of the criteria on the Filter toolbar, you must click the Filter button for those changes to take affect.

Insights Alerts Tab

Alerts send emails based on the thresholds that you have configured. These alerts enable you to be informed when exceptions occur in the CDN environment based on your desired thresholds.

You can set alerts based on the following:

- Bandwidth
- 4xx Errors
- 5xx Errors

- Edge Cache Hit Ratio

Metrics	Alert Condition *	Alert Name *	Email Id *	Frequency *	
Bandwidth	Equals To ▾ Enter value	Gbps Enter Alert	Enter Email	Every 15 Minutes ▾	Create
4xx Errors	Equals To ▾ Enter value	% Enter Alert	Enter Email	Every 15 Minutes ▾	Create
5xx Errors	Equals To ▾ Enter value	% Enter Alert	Enter Email	Every 15 Minutes ▾	Create
Edge Cache Hit Ratio	Equals To ▾ Enter value	% Enter Alert	Enter Email	Every 15 Minutes ▾	Create

Threshold Alerts					
	Metrics ▾	Alert Name ▾	Alert Condition ▾	Email Id ▾	Action ▾
1	4xx Errors	4xx Above 50 Percent	Grows Above 50	jsmith@company.com	Delete

Additionally, the following two alarms are not displayed on the **Alerts** tab, but will send email alerts when specific servers are unreachable:

- **Mid cache unreachable:** Sends an email whenever an Edge cache tries to pull content from a Mid cache and it cannot reach the Mid cache.
- **Origin Server unreachable:** Sends an email whenever a Mid cache tries to pull content from an Origin Server and it cannot reach the Origin Server.

Add an Alert



Note

If you are running OMD Director in an HA configuration, you can only make configuration changes from the Primary OMD Director instance. To perform the steps in this section, make sure you are logging into the Primary OMD Director instance.

Follow these steps to create a new alert:

1. From the Alerts tab, enter values for the following fields in the row of the metric you want to create the alert for:
 - For the Alert Condition, choose the condition to trigger the alert (Equals to, Grows Above, or Falls Below) and then enter a threshold value.
 - Enter a name for the Alert.
 - Enter an email address to which the alert should be sent. To add multiple email values, separate the email addresses by commas.
 - In the Frequency drop-down list choose the frequency for triggering the alert emails.

2. When you are finished entering the values for the alert, click **Create**. The alert will appear in the Threshold Alerts table.

Insights Custom Dashboards Tab

From the Custom Dashboards tab you can create and save custom dashboards based on custom searches that you enter.

Add a Custom Dashboard



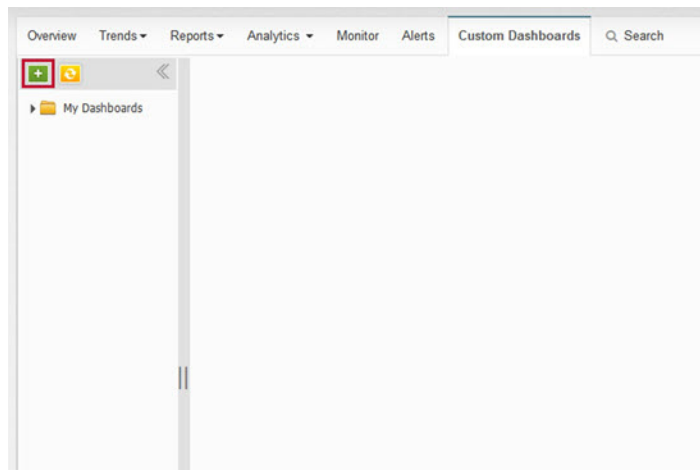
Note

If you are running OMD Director in an HA configuration, you can only make configuration changes from the Primary OMD Director instance. To perform the steps in this section, make sure you are logging into the Primary OMD Director instance.

Follow these steps to add a custom dashboard:

Procedure

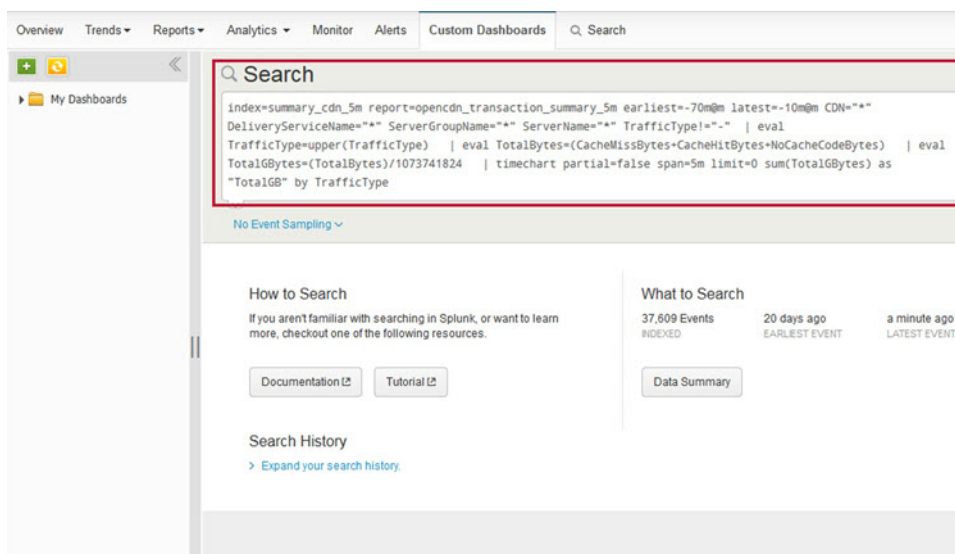
- Step 1** From within the Custom Dashboards tab, click the + in the upper left corner.



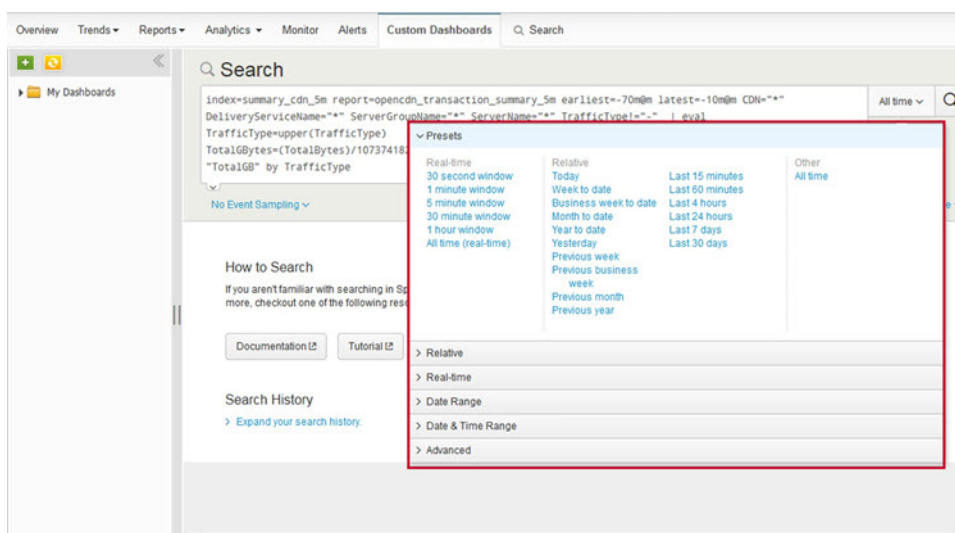
- Step 2** In the Search field that appears on the right, enter your query from which you want to create the dashboard.

Note

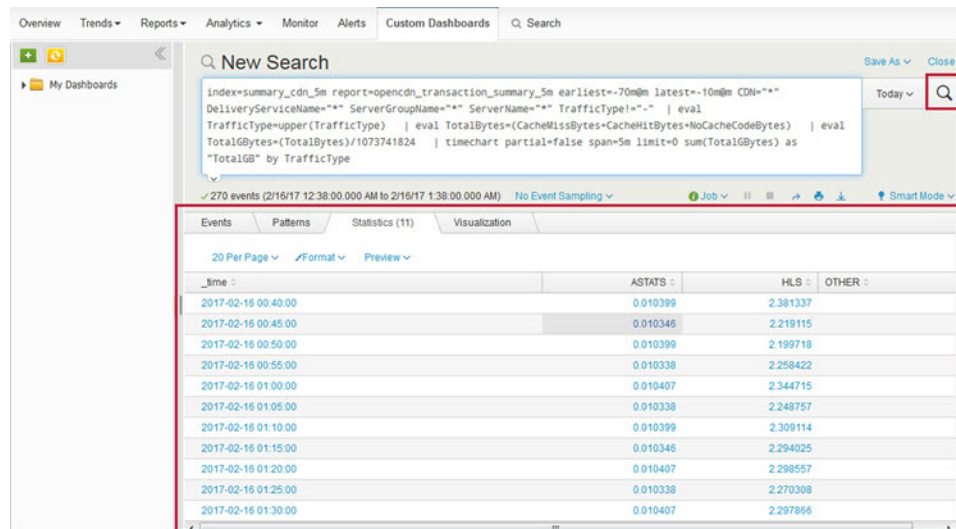
To change the formatting of numeric data to use a period to separate groups of thousands (instead of a comma), and to use commas for decimal places (instead of periods), enter “[converttoitalianformat” (without the quotes) at the end of any search query.



Step 3 From the Time drop-down list, choose the time frame that you would like to search across.



Step 4 Click the **Search** icon to confirm that the query works.



Step 5 From the **Save As** drop-down list, choose **Dashboard Panel**. The Save As Dashboard Panel window appears.

Step 6 In the Save As Dashboard Panel, to create a new dashboard, click **New** and configure the following information:

- **Dashboard Title (Optional):** Enter the custom dashboard title.

Note It is recommend that you always start the custom dashboard names with the word Custom (or any word other than dashboard). Using the word Custom helps the users to add panels to existing custom dashboard.

- **Dashboard ID:** Enter the custom dashboard ID.

Note The custom dashboard ID can contain only alphanumeric characters and underscores.

- **Dashboard Description (Optional):** Enter the custom dashboard description.

- **Dashboard Permissions:** Select the custom dashboard permissions as Private or Shared in App.

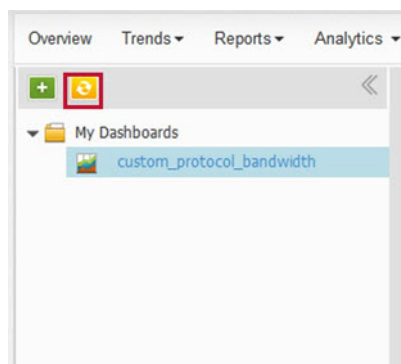
- **Private:** Only you can view and edit the custom dashboard.
- **Shared in App:** All users can view the created custom dashboard.

- **Panel Title (Optional):** Enter the panel title.

Step 7 Click **Save** to save the dashboard. Your Dashboard Panel has been Created dialog box appears.

Step 8 Click **View Dashboard** to view the newly created dashboard.

Note You may need to click the Refresh icon to see your newly created dashboard under My Dashboards in the navigation pane.



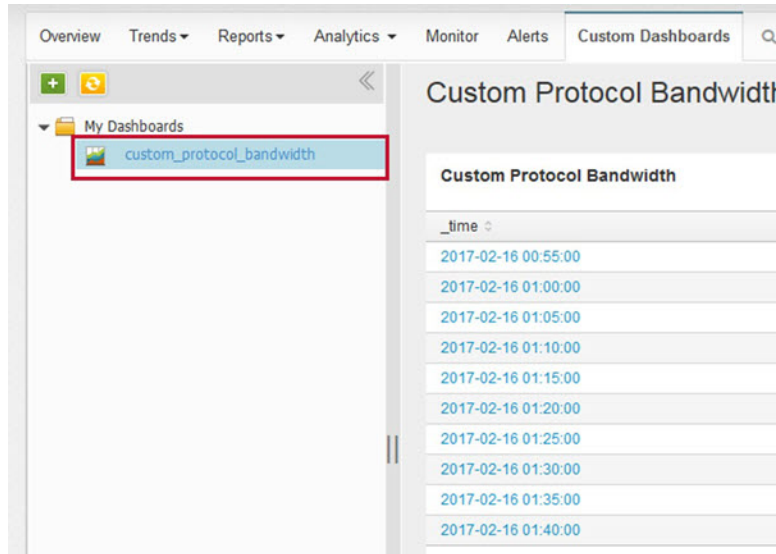
Delete a Custom Dashboard



Note If you are running OMD Director in an HA configuration, you can only make configuration changes from the Primary OMD Director instance. To perform the steps in this section, make sure you are logging into the Primary OMD Director instance.

Procedure

- Step 1** From the Custom Dashboard tab, under My Dashboards in the navigation panel, choose the custom dashboard that you want to delete.



- Step 2** From the Edit drop-down menu, choose **Delete**.
- Step 3** In the Delete confirmation dialog box that appears, click **Delete**.

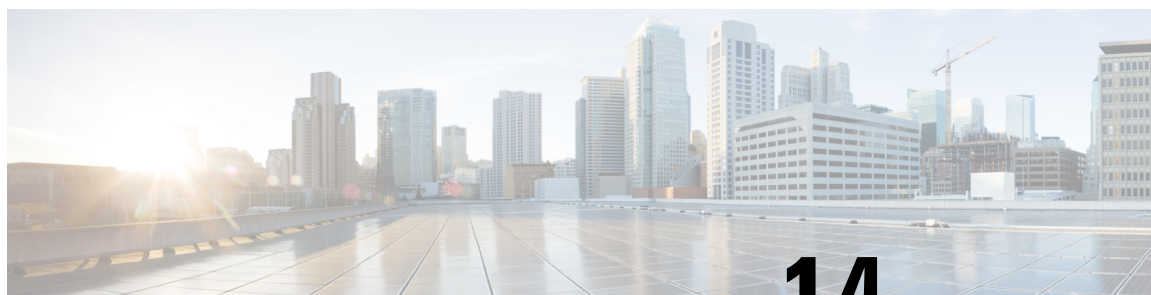
Insights Search Tab

All of the graphs, charts, and reports that you can access from the Overview, Trends, Reports, and Analytics tab are preconfigured as part of the OMD Insights installation. If you want to retrieve information that is not part of the predeveloped reports, you can create custom searches from the Search tab. The searches that you create and run from this tab are searching the analytics transactions database. These are the same type of searches from which you can build custom dashboards. In addition to entering the search criteria, you can also determine the time range across which the query will be executed.

From the Search tab window you can also see the total number of events that are currently in the transaction database. These are the total number of transactions that will be searched when you run a custom search.

For more information on the syntax of how to create a custom search, click the **Documentation** or **Tutorial** button from the Search tab.





CHAPTER 14

OMD Administration

The Administration menu enables you to perform the following administrative tasks:

- Manage the users who have access to OMD Director, including their privileges, and reset their passphrase.
- View a list of organizations that have been assigned OMD Director users and see which users have been assigned to these organizations.
- Manage the backups of the OMD databases, including adding remote backup destinations.
- Display an activity log that shows both user authentication and authorization activity, and CDN configuration activity.

This chapter describes how to perform these tasks in Cisco OMD Director.



Note If you are running OMD Director in an HA configuration, you can only make configuration changes from the Primary OMD Director instance. To perform the steps in this chapter, make sure you are logging into the Primary OMD Director instance. If you are logged into an OMD Director instance that is running in Backup mode, the header of the window will show “This OMD Director is currently running in backup mode”. If you are logged into an OMD Director running in a detached state because of an HA failure, the header will show “This OMD Director is currently running in detached mode”.



Note You must be logged in as a user with the Administrator role to access the Administration menu.

This chapter also describes how to change the user settings for your account in OMD Director, including how to change your passphrase, how to configure what alarms you would like to receive notifications for, and how you would like to receive those notifications.

This chapter includes the following topics:

- [Users/Organizations, on page 236](#)
- [Backup Management, on page 244](#)
- [User Profile, on page 251](#)
- [Activity Log, on page 255](#)

Users/Organizations

When you choose Administration > Users/Organizations from the navigation panel, the Users/Organizations page appears. This page has two sections:

- **Users:** Displays a list of users.
- **Organizations:** Displays a list of organizations and which users have been assigned to them.

Users

From the Users tab of the Users/Organizations page you can do the following:

- View user information
- Add a new user
- Edit an existing user
- Delete a user
- Reset a user passphrase

Figure 6: Users Tabs

Name	Username	Email Address	Phone Number	Expiration Date	
UserOne Administrator	User1	User1@company.example		0001-01-01 00:00	[edit] [delete] [lock?]
UserTwo CDN Operator (read & write)	User2	User2@company.example		0001-01-01 00:00	[edit] [delete] [lock?]
UserThree CDN Operator (read & write)	User3	User3@company.example		0001-01-01 00:00	[edit] [delete] [lock?]
UserFour Administrator	User4	User4@company.example		0001-01-01 00:00	[edit] [delete] [lock?]
UserFive Administrator	User5	User5@company.example		0001-01-01 00:00	[edit] [delete] [lock?]

Add a New User



Note

If you are running OMD Director in an HA configuration, you can only make configuration changes from the Primary OMD Director instance. To perform the steps in this section, make sure you are logging into the Primary OMD Director instance.

To create a new user, perform the following steps:

Procedure

- Step 1** From the Users tab of the Users/Organizations page, in the header of the right most column, click the blue + icon.

Name	Username	Email Address	Phone Number	Expiration Date	
UserOne Administrator	User1	User1@company.example		0001-01-01 00:00	[+]
UserTwo CDN Operator (read & write)	User2	User2@company.example		0001-01-01 00:00	[edit] [trash] [lock]
UserThree CDN Operator (read & write)	User3	User3@company.example		0001-01-01 00:00	[edit] [trash] [lock]
UserFour Administrator	User4	User4@company.example		0001-01-01 00:00	[edit] [trash] [lock]
UserFive Administrator	User5	User5@company.example		0001-01-01 00:00	[edit] [trash] [lock]

- Step 2** The Create New User window appears. Enter the following information for the new user:

- **Name:**

- **First:** Enter the first name for the user. This is a required field and can only contain letters, spaces, and dashes (-).
- **Middle:** Optionally enter a middle name. This field can only contain letters, spaces, and dashes (-).
- **Last:** Optionally enter a last name. This field can only contain letters, spaces, and dashes (-).

Note The user can change these fields.

- **Username:** Enter a username for the user. The username is what the user uses to log in to OMD Director. The username is case sensitive, must be at least 5 characters long, and can only contain letters and numbers. This is a required field.
- **Role:** Choose a role for the user. The following roles are available:
 - **Administrator:** Users assigned this role can access all features of OMD Director for all CDNs and CDN objects, including creating and managing users.
 - **CDN Admin:** Users assigned this role can perform all functions that are available from the Monitor, Provisioning, and Insights menu. CDN Admins cannot create or manage user accounts.
 - **CDN Viewer:** Users assigned this role can view information about the CDN and its objects, but they cannot change any of the objects or add any new objects. They also cannot see any information about the OMD Director users.
 - **Content Provider Admin:** Users assigned this role can only view Insights information, and only for Delivery Services that are assigned an organization that matches the Content Provider organization assigned to the user. Users with this role can also create custom dashboards in OMD Insights and perform custom searches. However, any custom dashboards that are created cannot be shared with other users.

- **Content Provider Viewer:** Users assigned this role can only view OMD Insights information, and only for Delivery Services that are assigned an organization that matches the Content Provider organization assigned to the user.
- **Reseller Admin:** Users assigned this role can only view OMD Insights information, and only for Delivery Services that are assigned an organization that is in the Reseller organization assigned to the user. Users with this role can also create custom dashboards in OMD Insights and perform custom searches. However, any custom dashboards that are created cannot be shared with other users.
- **Reseller Viewer:** Users assigned this role can only view OMD Insights information for Delivery Services that are assigned an organization that is in the Reseller organization assigned to the user.

Note You assign an organization to the Delivery Services using the Content Provider field on the Delivery Service.

- **Organization:** If you have assigned the Content Provider Admin, Content Provider Viewer, Reseller Admin, or Reseller Viewer role to the user, you must also assign an organization to the user. This organization determines the Delivery Services for which the user can view Insights information.

Note Reseller organizations can only be assigned to users with a "Reseller" role and Content Provider organizations can only be assigned to users with a "Content Provider" role. For more information on the different types of organizations, see the [Organizations](#) section.

Note You assign an organization to the Delivery Services using the Content Provider field on the Delivery Service.

- **Email Address:** Enter the email address for the client. This is a required field.
 - If you have assigned an SMTP server to be used with OMD Director, this must be a valid email address for the user. This email address is used to send the user a temporary passphrase for their initial login. When the user logs in for the first time using this temporary passphrase, they will be prompted to change their passphrase. This email address is also used for notifications that the user configures if the user chooses Email as a Notify Through option.
 - If you have not assigned an SMTP server to be used with OMD Director, you can enter any email address as long as it is in the format user@server.com. In this case, after you create the user, the temporary password for the user will display on the screen. You will need to provide this password to the user.

Note The user can change this field.

- **Phone Number:** Optionally enter a phone number for the user. This phone number is also used for notifications that the user configures if the user chooses SMS as a Notify Through option.

Note The user can change this field.

Step 3 Click **Add** to add the new user.

Step 4 If the user is created successfully, you will receive a confirmation message at the top of the Create New User window.

- If you have assigned an SMTP server to be used with OMD Director, the user will be sent an email with a temporary passphrase. This temporary passphrase is used the first time the user logs in to OMD Director.

Upon their first log in, they will be prompted to change this temporary passphrase. The temporary passphrase expires 72 hours after it was created.

- If you have not assigned an SMTP server to be used with OMD Director, the temporary password for the user will display on the screen. You will need to provide this password to the user.

Step 5 Click **Close** to close this window, or click **Add** and repeat Step 2 to Step 4 to add another user.

Edit an Existing User












Note If you are running OMD Director in an HA configuration, you can only make configuration changes from the Primary OMD Director instance. To perform the steps in this section, make sure you are logging into the Primary OMD Director instance.

To edit an exiting user, perform the following steps:

Procedure

Step 1 From the Users tab of the Users/Organizations page, click the edit icon in the row of the user you want to edit.

Users					
Search User <input type="text"/>					
Name	Username	Email Address	Phone Number	Expiration Date	
UserOne Administrator	User1	User1@company.example		0001-01-01 00:00	  
UserTwo CDN Operator (read & write)	User2	User2@company.example		0001-01-01 00:00	  
UserThree CDN Operator (read & write)	User3	User3@company.example		0001-01-01 00:00	  

Step 2 The Edit User window appears. Update the desired fields. You can change the following fields for the user:

Note You cannot change the role or organization of a user that has the Content Provider Admin, Content Provider Viewer, Reseller Admin, or Reseller Viewer role. If you need to change either of these values for one of these users, you must delete the user and recreate them.

Step 3 When you are finished making your changes, click **Save**. If you want to exit the Edit User window without saving the changes, click **Close**.

Step 4 After you save the changes, a confirmation message appears at the top of the Edit User window. Click **Close** to close this window.

Delete a User



Note If you are running OMD Director in an HA configuration, you can only make configuration changes from the Primary OMD Director instance. To perform the steps in this section, make sure you are logging into the Primary OMD Director instance.

To delete a user, perform the following steps:

Procedure

Step 1 From the Users tab of the Users/Organizations page, click the **Delete** icon in the row of the user you want to delete.

Name	Username	Email Address	Phone Number	Expiration Date	
UserOne Administrator	User1	User1@company.example		0001-01-01 00:00	
UserTwo CDN Operator (read & write)	User2	User2@company.example		0001-01-01 00:00	
UserThree CDN Operator (read & write)	User3	User3@company.example		0001-01-01 00:00	

Step 2 The Delete User confirmation window appears. To finish deleting the user, click **Yes**.

Reset a User Passphrase












Note If you are running OMD Director in an HA configuration, you can only make configuration changes from the Primary OMD Director instance. To perform the steps in this section, make sure you are logging into the Primary OMD Director instance.

To reset the passphrase for a user, perform the following steps:

Procedure

Step 1 From the Users tab of the Users/Organizations page, click the reset passphrase icon (the lock with the question mark) in the row of the user whose passphrase you need to reset. Clicking this link will reset the user passphrase to a temporary passphrase.

Users					
Search User					
Name	Username	Email Address	Phone Number	Expiration Date	
UserOne Administrator	User1	User1@company.example		0001-01-01 00:00	  
UserTwo CDN Operator (read & write)	User2	User2@company.example		0001-01-01 00:00	  
UserThree CDN Operator (read & write)	User3	User3@company.example		0001-01-01 00:00	  

Step 2 After you click the reset passphrase icon:

- If you have assigned an SMTP server to be used with OMD Director, you will receive the message “Passphrase has been emailed successfully for *username*”. This indicates that the user passphrase has been successfully reset to a temporary passphrase and an email has been sent to the user. The email will direct the user to log in to OMD Director to change their passphrase. They will have 72 hours to do this before the temporary passphrase expires.
- If you have not assigned an SMTP server to be used with OMD Director, you will receive the message "Reset Passphrase Successfully" and the new temporary password will display on the screen. You will need to provide this password to the user.

Organizations

From the Organizations tab of the Users/Organizations page you can do the following:

- View organizations
- Add organizations
- Delete organizations

Users that are assigned one of the following roles must be assigned an organization. The assigned organization determines the Delivery Services for which they can view analytics information:

- Content Provider Admin
- Content Provider Viewer
- Reseller Admin
- Reseller Viewer

Reseller users can only view Insights information for Delivery Services that are assigned an organization that belongs to the Reseller organization assigned to the user. Content Provider users can only view Insights information for Delivery Services that are assigned an organization that matches the Content Provider organization they are assigned. Also, users that are assigned a Content Provider or Reseller role can only access the Insights menu and the Support > Contact menu in OMD Director. They will not have access to any other menus in OMD Director.

There are two types of organizations:

- **Reseller:** Reseller organizations contain Content Provider organizations, which allows an organizational hierarchy to be created. Reseller organizations can only be assigned to users with a "Reseller" role. Reseller users can only view Insights information for Delivery Services that are assigned an organization that belongs to the Reseller organization assigned to the user.
- **Content Provider:** Content Provider organizations can be assigned to a Reseller organization *or* can be a "global" (standalone) organization. Content Provider organizations can only be assigned to users with a "Content Provider" role.
 - If a Content Provider organization is assigned to a Reseller organization, the following users can view the analytics for Delivery Services assigned this organization:
 - Content Provider users whose organization is set to this specific Content Provider organization
 - Reseller users whose organization is set to a Reseller organization to which this Content Provider organization has been assigned
 - If a Content Provider organization is "global", only Content Provider users that have been assigned this specific organization can view the analytics for the Delivery Services assigned this organization. No Reseller users would be able to view the analytics for Delivery Services assigned this organization because the Content Provider organization was *not* assigned to a Reseller organization.



Note Only Content Provider organizations are assigned to Delivery Services.

View Organizations

The Organizations tab of the Users/Organizations window displays a list of organizations and which users have been assigned to each organization. If an organization has been assigned to a user, this organization determines the Delivery Services for which the user can view Insights information. The user can only view information for Delivery Services that are assigned an organization that matches the Content Provider organization assigned to the user or that is in the Reseller organization assigned to the user.



Note Organizations are only assigned to users that have the Content Provider Admin, Content Provider Viewer, Reseller Admin, or Reseller Viewer role.

Organizations		
		Search Organ <input type="text"/>
Name	Users	
cp-two	cp2-viewer-1	
reseller-one	reseller-admin-1	
cp-one	cp1-viewer-1	

Add a New Organization



Note If you are running OMD Director in an HA configuration, you can only make configuration changes from the Primary OMD Director instance. To perform the steps in this section, make sure you are logging into the Primary OMD Director instance.

To create a new organization, perform the following steps:

Procedure

Step 1 From the Organizations tab of the Users/Organizations page, in the header of the right most column, click the **Create New Organization (+)** icon.

Organizations		
		Search Organ <input type="text"/>
Name	Users	
cp-two	cp2-viewer-1	
reseller-one	reseller-admin-1	
cp-one	cp1-viewer-1	

Step 2 The Create New Organization window appears.

Create New Organization

Organization Name *

Organization Type *

Content Provider

Resellers

Reseller-One

Close

Add

Enter the following information for the new organization:

- **Organization Name:** Enter a name for the organization.
- **Organization Type:** Choose either Reseller or Content Provider.
- **Resellers:** If the Organization Type that you chose was Content Provider, choose optionally choose the Reseller that the Content Provider organization belongs to.

Note You cannot edit an organization after you create it. If you need to change the organization type or Reseller organization that a Content Provider organization is assigned to, you will need to delete the organization and recreate it.

- Step 3** Click **Add** to add the new organization.
- Step 4** If the organization is created successfully, you will receive a confirmation message at the top of the Create New Organization window.
- Step 5** Click **Close** to close this window, or repeat Step 2 and Step 3 to add another user.

Delete an Organization and its Users

If you are deleting an organization that has users in it, you will be prompted to also delete the users. You cannot delete the organization without also deleting the users that are assigned to it.



Note If you are running OMD Director in an HA configuration, you can only make configuration changes from the Primary OMD Director instance. To perform the steps in this section, make sure you are logging into the Primary OMD Director instance.

Follow these steps to delete an organization and the users assigned to it:

Procedure

- Step 1** In the row for the organization that you want to delete, click the **Delete** icon.

Organizations		
		Search Organ <input type="text"/>
Name	Users	
cp-two	cp2-viewer-1	
reseller-one	reseller-admin-1	
cp-one	cp1-viewer-1	

- Step 2** The Delete Organization window appears asking you to confirm the deletion of the organization and any users that are associated with that organization. To confirm the deletion, click **Yes**. To cancel the deletion of the organization and its users, click **No**.

Backup Management



Note If you are running OMD Director in an HA configuration, you can only make configuration changes from the Primary OMD Director instance. To perform the steps in this section, make sure you are logging into the Primary OMD Director instance.

When you install OMD Director, it automatically schedules local daily backups of the databases that are managed by the OMD Director Worker node. These backups are saved in the

/home/ocdn_adm/director/backup_ms/backup_files/ folder on the OMD Director Worker node. You can also configure OMD Director to copy these backups to a remote server, which is recommended for disaster recovery, and you can change the frequency and day on which the backups occur. To configure these items and additional backup management features, go to **Administration > Backup Management**.

From the Backup Management window in OMD Director you can do the following:

- View a list of scheduled backups
- Change the frequency and time of a scheduled backup
- Enable or disable a scheduled backup
- Perform a backup or restore of a database
- Add remote backup destinations
- Modify remote backup destinations
- Add additional databases to back up
- Modify the settings for the databases that are already configured for back up

This section will cover how to perform these functions.

Schedule Backups

By default OMD Director is configured to perform a daily backup of the following OMD Director databases:

- alarms_db
- provision_db (also referred to as the CDN Manager database)
- user_management
- omd_notifications

You can see a list of the scheduled backups and make changes to when they occur from the Schedule page. To access the Schedule page, from the Backup Management window click the **Schedule Backup** tab. The Schedule page appears showing a list of databases to be backed up, the day and time they will be backed up, and whether the scheduled backup is enabled.

From the Schedule page you can do the following:

- Change the frequency and time of the backup for each database
- Enable or disable the scheduled backup
- Manually back up of a database
- Restore a database

Backup Management							
Schedule Backup							
Schedule							
Database Name	Database Type	Schedule				Last Run Time	Last Run Status
		Day	Hour	Minute	second		
alarms_db	mongo	Every Day	00	05	00	Wed 13 Sep 2017 05:05:00 PM	success
provision_db	mongo	Every Day	00	10	00	Wed 13 Sep 2017 05:10:00 PM	success
user_management	postgres	Every Day	00	15	00	Wed 13 Sep 2017 05:15:00 PM	success
omd_notifications	postgres	Every Day	00	20	00	Wed 13 Sep 2017 05:20:00 PM	success

Change Schedule

The following is the default schedule for each database:

- **alarms_db:** Every day at 5 minutes after midnight
- **provision_db:** Every day at 10 minutes after midnight
- **user_management:** Every day at 15 minutes after midnight
- **omd_notifications:** Every day at 20 minutes after midnight

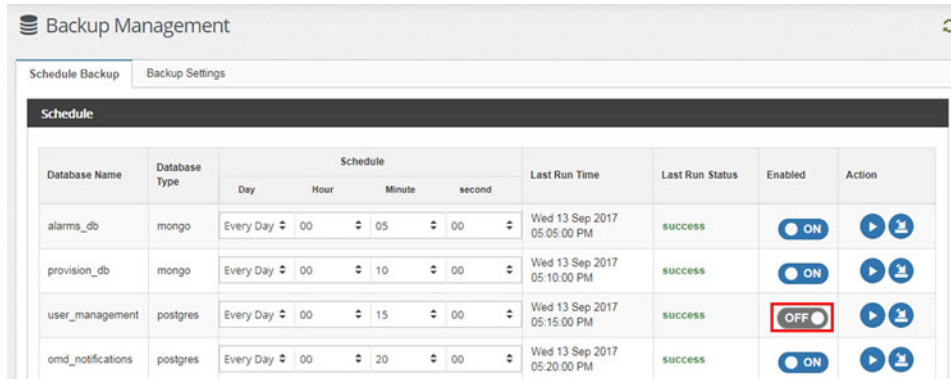
To back up the database only once a week instead of daily, from the Day drop-down list, choose the day of the week to back up the database. To change the time of day that the backup should occur, choose the new time that you want from the Hour, Minute, and Second drop-down lists as appropriate.

Schedule Backup							
Backup Settings							
Schedule							
Database Name	Database Type	Schedule					
		Day	Hour	Minute	second		
alarms_db	mongo	Every Day	00	05	00		
provision_db	mongo	Every Day	00	10	00		
user_management	postgres	Every Day	00	15	00		
omd_notifications	postgres	Every Day	00	20	00		
backups_test_db	mongo	Every Day	00	00	00		

Enable/Disable a Scheduled Backup

To disable the scheduled backup of a database, in the Enabled column of the database, slide the button to the right until you see “OFF”. To enable the scheduled backup, slide the button to the left until you see “ON”.

By default, the automatically created scheduled backups are enabled.



Backup Management									
Schedule Backup									
Schedule									
Database Name	Database Type	Schedule				Last Run Time	Last Run Status	Enabled	Action
		Day	Hour	Minute	second				
alarms_db	mongo	Every Day	00	05	00	Wed 13 Sep 2017 05:05:00 PM	success	ON	
provision_db	mongo	Every Day	00	10	00	Wed 13 Sep 2017 05:10:00 PM	success	ON	
user_management	postgres	Every Day	00	15	00	Wed 13 Sep 2017 05:15:00 PM	success	OFF	
omd_notifications	postgres	Every Day	00	20	00	Wed 13 Sep 2017 05:20:00 PM	success	ON	

Manually Back Up and Restore a Database

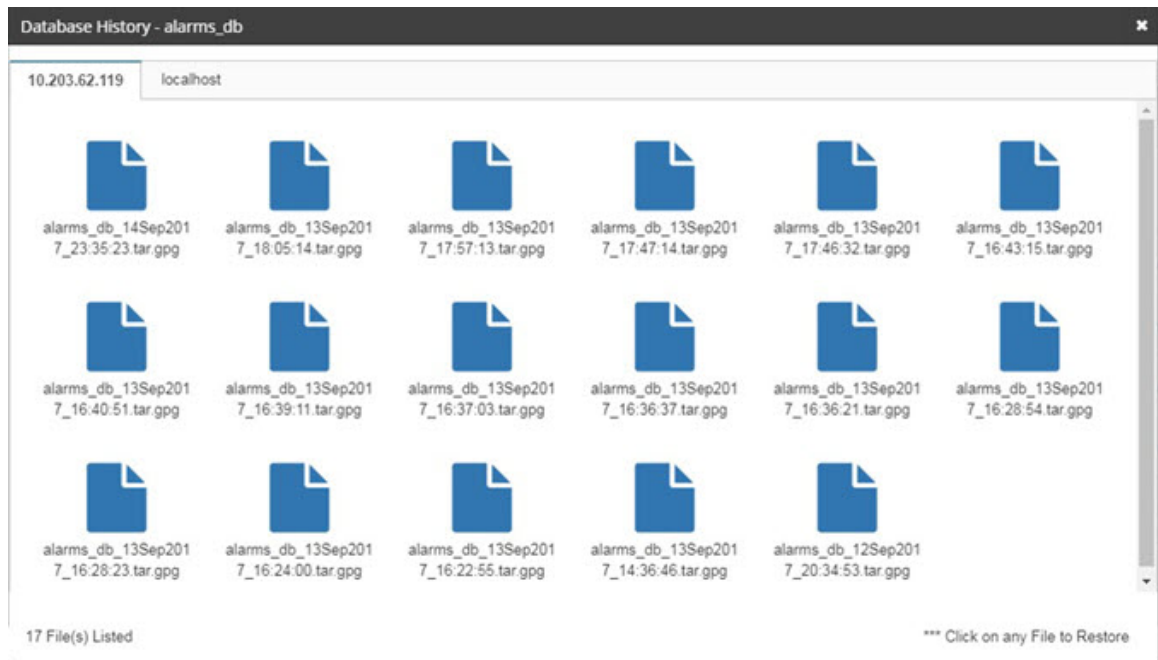
To perform a manual back up of a database, in the Action column of the database, click the **Backup Now** icon. If the backup is successful, you will receive a backup completed message. This will create a copy of the backup locally on the OMD Director Worker node and on each remote server that has been configured. By default, there are no remote servers configured.

To restore a database from a backup file, perform the following steps:

Procedure

Step 1 In the Actions column of the database, click the **Restore** icon. The Database History window appears. This window will have a localhost tab and a tab for every remote destination that is configured.

Note For the local copy of the backup, only the most recent successful backup is saved, so you will only have one file to choose on the localhost tab. For the remote copies of the backup, each successful backup that is run for a database will create a new file, so you will have multiple versions to choose from. The filenames will contain the date and time that the file was created to help you choose.



Step 2 Click the tab for the location that contains the backup file that you want to restore and then click the file to restore.

Step 3 When prompted, click **Yes** to confirm that you want to restore the database. If the restore is successful, you will receive a restore successful message.

Backup Settings

The Backup Settings tab contains two sections:

- **Backup Destinations:** By default the OMD Director databases are backed up locally to the `/home/ocdn_adm/director/backup_ms/backup_files/` folder on the OMD Director Worker node. It is also recommended that you configure OMD Director to copy these backups to a remote server for disaster recovery. From the Backup Destinations section you can add remote servers to copy the backup files to and you can test the connection to these servers.
- **Database Configuration:** From the Database Configuration section of the Backup Settings you can see information about the OMD Director databases that are currently being backed up and you can edit the settings for these databases. From this section you can also add additional Mongo or Postgres databases to back up. (The database must be on the OMD Director Worker node or a container on the OMD Director Worker node and you will need to provide the username and password for access.) Any databases that appear in the Database Configuration section will have a scheduled backup listed on the Schedule Backup tab.

Add Remote Backup Destinations

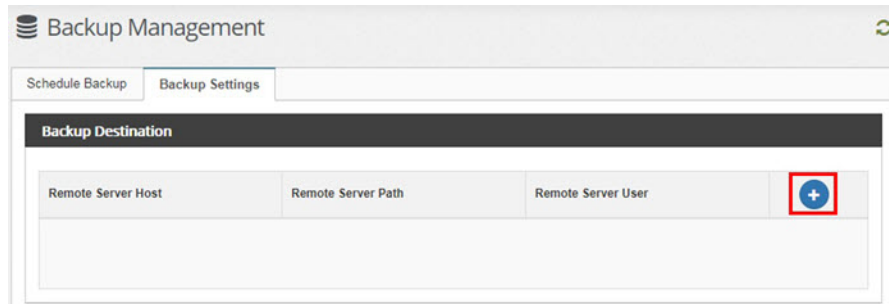
When you add a remote backup destination, every scheduled backup that is enabled in the Schedule list will save a copy of the backup file to the remote server, in addition to the local copy. For the local copy of the backup, only the most recent successful backup is saved. For the remote copies of the backup, each successful

backup that is run will create a new file; the filename will contain the date and time that the file was created. These files will be saved until you delete them. If you add multiple remote backup destinations, each destination will receive a copy of the database backup file for every scheduled backup.

To add a remote server as a database backup destination, perform the following steps:

Procedure

- Step 1** From the Backup Settings tab in the Backup Destination section, click the Create New Destination (+) icon in the column heading of the last column.



- Step 2** The Create Destination window appears. In this window, enter the values for the field as appropriate. The following table provides a description of the fields.

Table 3: Create Destination Fields

Fields	Description
Remote Server Host	Enter the IP address or hostname of the remote server
Remote Server Path	Enter the path on the remote server in which you want to save the database backups.
Remote Server User	Enter the username of a user that can SCP data to the server and that has access privileges to the Remote Server Path that is specified.
Private Key File	Browse for the private key file of the user you entered in the Remote Server User field.
Private Key Passcode	If the private key file requires a passcode, enter it in this field.

- Step 3** Click **Add** to save the settings and add the new destination. You should receive a message that the destination was created successfully.
- Step 4** After the remote server is added, click the **Test Connection** icon to confirm that you can connect to the server. You should receive the message that the destination was successfully verified.

Edit or Delete a Remote Backup Destination

To delete a remote backup destination, click the **Delete** icon for that destination. When prompted, click **Yes** to confirm the deletion of the backup destination.

To Edit the settings of a remote backup destination, click the **Edit** icon for that destination. The Edit Backup Destination page appears. Make the necessary changes to the settings, referring to the table in [Add Remote Backup Destinations, on page 248](#) for a description of the fields. When you are finished making changes, click **Save**.

Add a New Database Backup

To add a new Mongo or Postgres database to back up, perform the following steps:







Note

You can only add databases that are on the OMD Director Worker node or a container on the OMD Director Worker node.

Procedure

Step 1

From the Backup Settings tab in the Database Configuration section, click the Create New Database Backup (+) icon in the column heading of the last column.

Database Configuration			
Name	Database Type	Host	
alarms_db	mongo	10.254.141.213	
provision_db	mongo	10.254.169.231	
user_management	postgres	10.254.50.16	
omd_notifications	postgres	10.254.182.204	

Step 2

The Create New Database Backup window appears. In this window, enter the values for the field as appropriate. The following table provides a description of the fields:

Table 4: Create New Database Backup Fields

Fields	Description
Database Name	Enter the name of the database file to back up.
Database Type	Choose the type of database that is being backed up. The options are Mongo or Postgres.
Authentication Database	If the username that is used to log in to the database is not in the admin database, enter the authentication database for the user. Note: This field is only for Mongo databases.
Host	Enter the IP address of the server that contains the database. This can be either the OMD Director Worker node or a container on the OMD Director Worker node.
Port	Enter the port number used to connect to the database. By default Mongo databases use 27017 and Postgres databases use 5432.

Fields	Description
Username	Enter the username needed to log in to the database.
Password	Enter the password for the username that can log in to the database.

- Step 3** Click **Add** to save the settings and add the new database configuration. You should receive a message that the database backup was created successfully.
- Step 4** After the database configuration is added, an entry is automatically added to the Schedule list on the Schedule Backup tab. By default, the scheduled backup will be configured to run daily at midnight and will be disabled. Refer to [Schedule Backups, on page 245](#) to change these settings.

Edit a Database Configuration

To Edit the settings of a database configuration, click the **Edit** icon for that database. The Edit Database Configuration page appears. Make the necessary changes to the settings, referring to [Add a New Database Backup, on page 250](#) for a description of the fields. When you are finished making changes, click **Save**.

User Profile



Note If you are running OMD Director in an HA configuration, you can only make configuration changes from the Primary OMD Director instance. To perform the steps in this section, make sure you are logging into the Primary OMD Director instance.

When you are logged into OMD Director, from the Profile page you can do the following:

- Change the following information:
 - Your first name, middle name, and last name
 - Your email address
 - Your phone number
 - Your passphrase. Your passphrase must meet the following criteria:
 - By default it must be at least 8 characters but should not exceed 47 characters. The default minimum length is configurable.
 - By default it must contain:

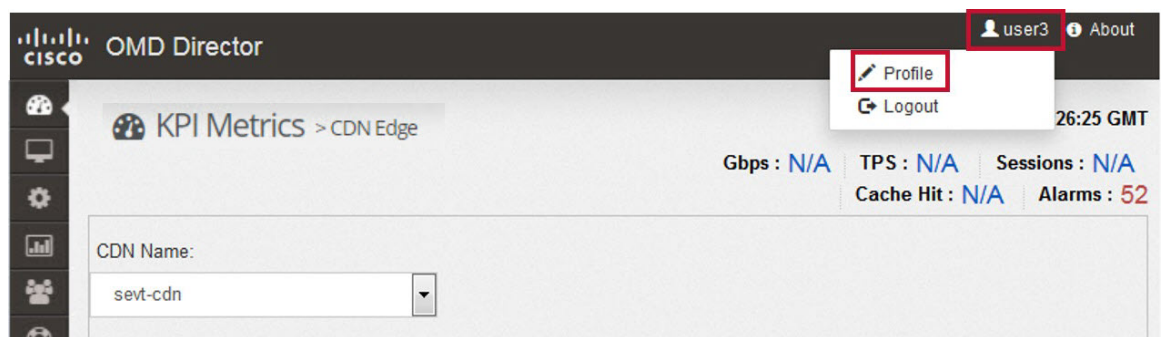


Note Which categories are required is based on the configuration of the OMD Director password parameters in the omd.conf file.

- At least one letters
- At least one number

- At least one special character
- It cannot contain €, @, #, " or ,
- By default it cannot contain spaces. This is configurable by the system administrator in the omd.conf file.
- Configure your notification settings.

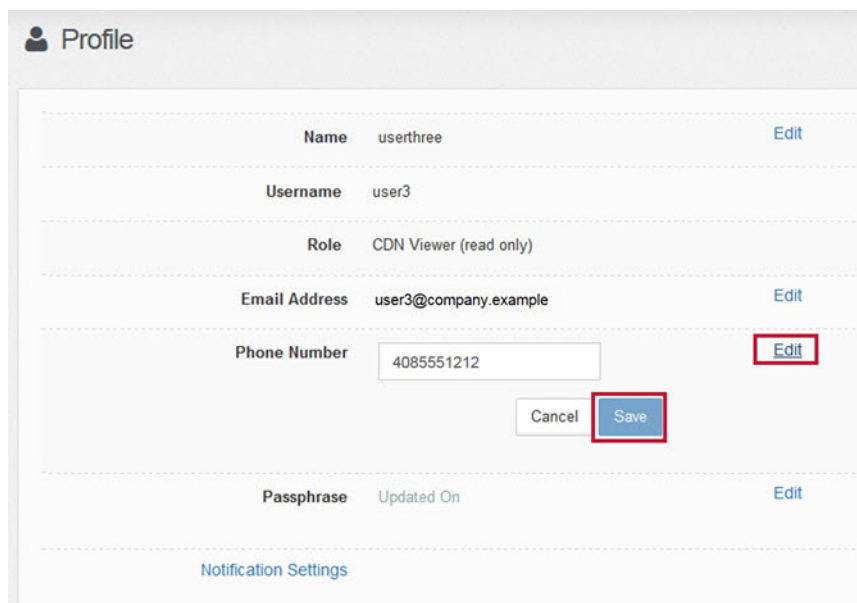
To access the Profile page to make these changes, on the upper right side of the OMD Director window, click your username and choose **Profile**.



The Profile page appears.

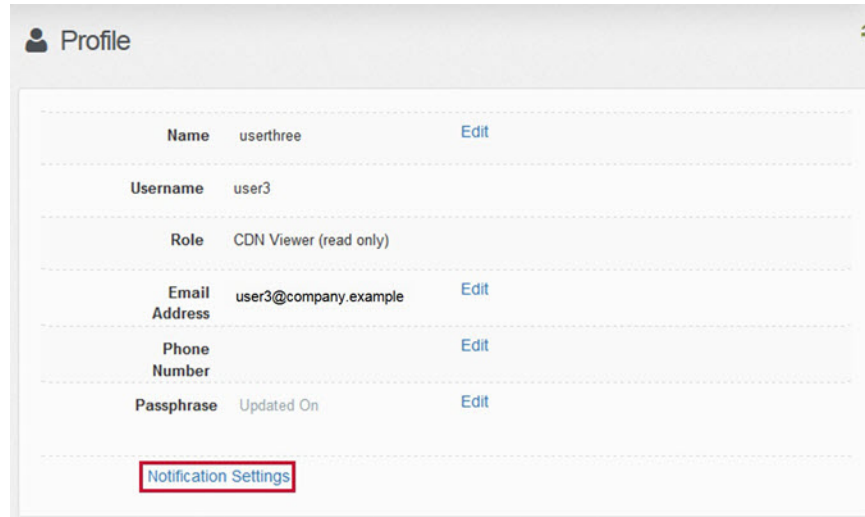
Change User Settings

To change your name, email address, or phone number click **Edit** next to the item that you want to change on the Profile page. When you are done making your change click **Save** to save the change.



Notification Settings

From the Profile page you can configure your alarm notification settings. To access these settings, at the bottom of the Profile page, click the **Notification Settings** link.



The screenshot shows a 'Profile' page with a header bar containing a user icon and the title 'Profile'. Below the header is a table of user information. At the bottom of the table, there is a link labeled 'Notification Settings' which is highlighted with a red rectangular box.

Name	userthree	Edit
Username	user3	
Role	CDN Viewer (read only)	
Email Address	user3@company.example	Edit
Phone Number		Edit
Passphrase	Updated On	Edit
Notification Settings		

The Notification Settings page appears. From this page you can configure which alarm names you would like to receive notifications for and what criteria for those alarms should trigger a notification. The criteria can be:

- If a new alarm is generated
- If the severity of an alarm has changed
- If an alarm is acknowledged
- If an alarm is cleared
- If an alarm is resolved

Profile > Notification Settings « back

Notification Settings

Notify ☒ Email ☐ SMS Preset Save

Through*

Alarm Name	Status					
	All	New	Severity Changed	Acknowledged	Cleared	Resolved
cpu_usage	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
disk_usage	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
dns_status	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
keep_alive	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
ntp_status	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
password_failed_attempt	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
ram_usage	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
service_sshd_status	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
service_tripwire_status	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
tcp_connection	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Preset Save

To choose which alarm names and which criteria for those alarms you would like to receive notifications for, you have the following options:

- Check the criteria check box for the individual alarm names for which you would like to generate notifications.
- To receive notifications for all of the alarm names for a specific criteria, check the check box in the column header for that criteria. For example, if you would like to receive a notification for any new alarms that are generated, regardless of what alarm name they are generated for, click the check box in the New column header.
- To receive notifications for all of the criteria for a specific alarm name, check the check box in the All column for that alarm name.
- If you would like to receive notifications for all of the criteria for all of the alarms, check the check box in the All column header. You can also use this check box to choose all of the alarms and then deselect the criteria for the alarms for which you do not want to receive notifications.

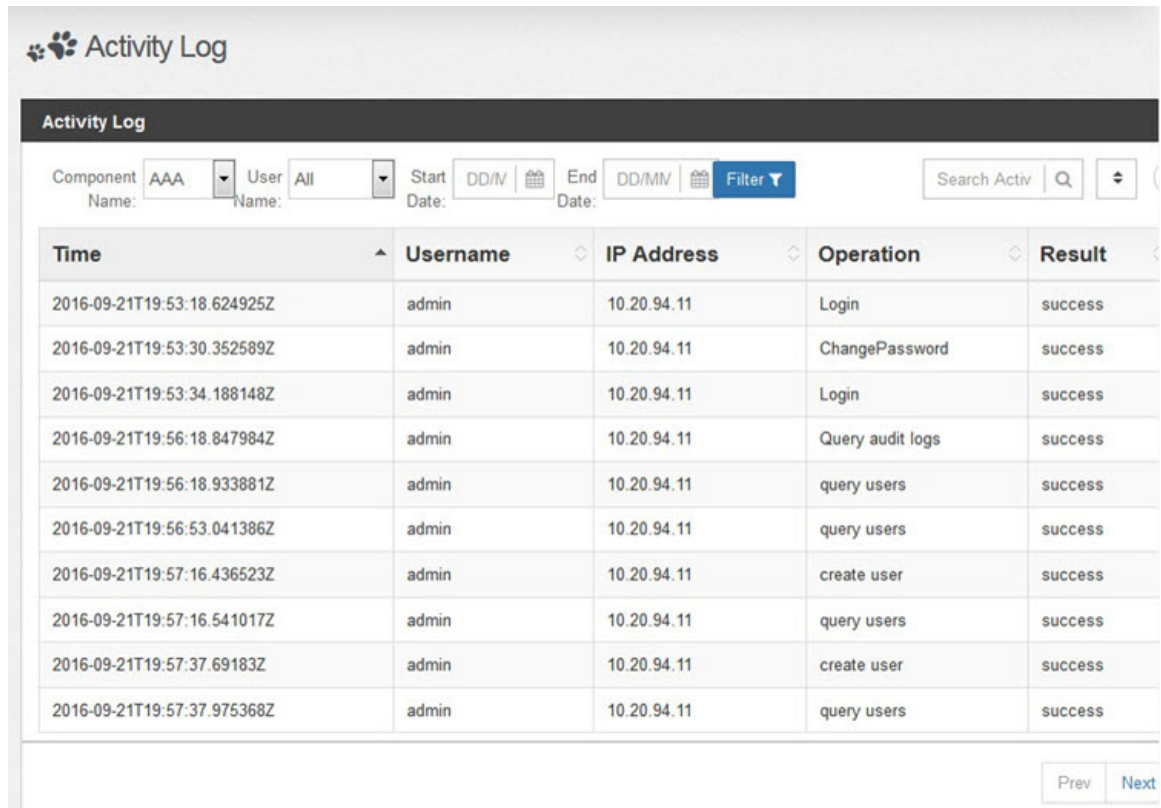
To actually receive the notifications, you must also configure whether you would like the notifications sent to you via email. To receive notifications via email, check the **Email** check box at the top left of the Notification Settings page. Email notifications are sent to the email address that is configured for your user account in OMD Director. If you need to change this setting, see [Change User Settings, on page 252](#).



Note The SMS option is not supported in the current release.

Activity Log

The Activity Log enables you to view an activity log that shows both user authentication and authorization activity, and CDN configuration activity. To access the Activity Log, choose **Administration > Activity Log** from the navigation panel.



Time	Username	IP Address	Operation	Result
2016-09-21T19:53:18.624925Z	admin	10.20.94.11	Login	success
2016-09-21T19:53:30.352589Z	admin	10.20.94.11	ChangePassword	success
2016-09-21T19:53:34.188148Z	admin	10.20.94.11	Login	success
2016-09-21T19:56:18.847984Z	admin	10.20.94.11	Query audit logs	success
2016-09-21T19:56:18.933881Z	admin	10.20.94.11	query users	success
2016-09-21T19:56:53.041386Z	admin	10.20.94.11	query users	success
2016-09-21T19:57:16.436523Z	admin	10.20.94.11	create user	success
2016-09-21T19:57:16.541017Z	admin	10.20.94.11	query users	success
2016-09-21T19:57:37.69183Z	admin	10.20.94.11	create user	success
2016-09-21T19:57:37.975368Z	admin	10.20.94.11	query users	success

The default page that appears displays the user authentication and authorization activity. From this page you can do the following for the AAA log entries:

- **Filter the information by user and date:** To filter the log entries based on these criteria, choose the user from the User drop-down list and optionally choose a Start Date and End Date. When you are done entering the filter criteria, click the **Filter** button.
- **Search the log entries:** To search the entries, enter the data you want to search for in the Search box. The text you enter is searched across all of the fields in the log entries.
- **Choose how many rows to display at a time:** You can choose this from the field next to the Search box. The options are 5, 10, 25, and 50. The default is 10.
- **Save the activity log as a JSON or CSV file:** To save the AAA activity to a file, click the download arrow at the end of the toolbar and choose the file format to save.

To view the CDN configuration activity, choose **CDN Configuration** from the Component Name drop-down list in the toolbar and click the **Filter** button.

Activity Log

Component Name: CDN Conf Filter Search Activity Q ⌵ ⌵

Time	Username	IP Address	Operation	Result
Thursday, March 02, 2017 13:28:03 UTC	user2	10.20.63.9	ServerHandler PUT	SUCCESS
Thursday, March 02, 2017 13:27:04 UTC	user2	10.20.63.9	ServerHandler PUT	SUCCESS
Wednesday, March 01, 2017 19:41:37 UTC	user2	10.20.63.9	CacheGroupHandler DELETE	SUCCESS
Wednesday, March 01, 2017 19:39:54 UTC	user2	10.20.63.9	CacheGroupHandler PUT	SUCCESS
Wednesday, March 01, 2017 19:39:30 UTC	user2	10.20.63.9	CacheGroupHandler PUT	SUCCESS
Wednesday, March 01, 2017 00:28:14 UTC	user2	10.20.63.9	CacheGroupHandler POST	FAILED
Tuesday, February 28, 2017 15:24:55 UTC	user2	10.20.63.9	DeliveryServiceHandler PUT	FAILED
Tuesday, February 28, 2017 15:24:53 UTC	user2	10.20.63.9	DeliveryServiceHandler PUT	FAILED
Tuesday, February 28, 2017 14:52:50 UTC	user2	10.20.63.9	CacheGroupHandler PUT	SUCCESS
Tuesday, February 28, 2017 14:47:49 UTC	user2	10.20.63.9	DeliveryServiceHandler PUT	SUCCESS

From this page you can do the following for the CDN configuration log entries:

- **Search the log entries:** To search the entries, enter the data you want to search for in the Search box. The text you enter is searched across all of the fields in the log entries.
- **Choose how many rows to display at a time:** You can choose this from the field next to the Search box. The options are 5, 10, 25, and 50. The default is 10.
- **Save the activity log as a JSON or CSV file:** To save the AAA activity to a file, click the download arrow at the end of the toolbar and choose the file format to save.



APPENDIX **A**

Example CZF File

The CZF file is a static JSON file that is used by the Traffic Routers to help determine which clients are directed to which cache groups. The CZF file specifies which client IP addresses to redirect to which cache groups. If the location (longitude and latitude) of the cache groups are configured in the CZF file and the “Fallback Enable” setting of the mapped Edge cache group is set to “Yes”, when no caches are available in the first Cache Group matched, the Traffic Router will use the longitude and latitude to find the next closest cache group to use as a backup. For more information on Backup Edge Groups, see [Backup Edge Cache Groups, on page 78](#).

Also if the Geo Limit setting for a Delivery Service is set to “CZF only”, the Traffic Router will use the CZF file to determine which client requests to allow. In this case, the Traffic Router will only allow requests from clients whose IP addresses match an entry in the CZF file.

The following is an example of the format of the CZF file:

```
{
  "coverageZones":
  {
    "Edge-West1":
    {
      "network6":
      [
        "1234:5704::/64",
        "1234:5705::/64",
        "1234:5706::/64"
      ],
      "network":
      [
        "192.168.4.0/24",
        "192.168.5.0/24",
        "192.168.6.0/24",
        "192.168.7.0/24",
        "192.168.8.0/24",
        "192.168.9.0/24"
      ],
      "coordinates":
      {
        "longitude": -118,
        "latitude": 34
      }
    },
    "Edge-West2":
    {
      "network6":
      [
        "1234:5710::/64",
```

```

        "1234:5711::/64",
        "1234:5712::/64"
    ],
    "network":
    [
        "192.168.10.0/24",
        "192.168.11.0/24",
        "192.168.12.0/24",
        "192.168.13.0/24",
        "192.168.14.0/24",
        "192.168.15.0/24"
    ],
    "coordinates":
    {
        "longitude": -122,
        "latitude": 47
    }
},
"Edge-West3":
{
    "network6":
    [
        "1234:5720::/64",
        "1234:5721::/64",
        "1234:5722::/64"
    ],
    "network":
    [
        "192.168.20.0/24",
        "192.168.21.0/24",
        "192.168.22.0/24",
        "192.168.23.0/24",
        "192.168.24.0/24",
        "192.168.25.0/24"
    ],
    "coordinates":
    {
        "longitude": -117,
        "latitude": 32
    }
},
"Edge-West4":
{
    "network6":
    [
        "1234:5730::/64",
        "1234:5731::/64",
        "1234:5732::/64"
    ],
    "network":
    [
        "192.168.30.0/24",
        "192.168.31.0/24",
        "192.168.32.0/24",
        "192.168.33.0/24",
        "192.168.34.0/24",
        "192.168.35.0/24"
    ],
    "coordinates":
    {
        "longitude": -105,
        "latitude": 40
    }
}
}

```

```
}  
}
```




APPENDIX B

Example Ingest Manifest File

The ingest manifest file is a JSON file that identifies what content should be prepositioned for a delivery service. This file can reference single files, directories, and ABRs. It can also contain optional attributes that control when to download the items, and how to protect them in cache. After the manifest is downloaded, the preposition daemon will parse it and download all segment files.

The "Ingest Manifest File Format" table provides information on the format of the ingest manifest file. Two example ingest manifest files are shown after the table.

Table 5: Ingest Manifest File Format

Key1	Key2	Key3	Description
cookie			(Optional) Customized cookie in request sent to Origin Server
headers			(Optional) Customized headers in request sent to Origin Server
preload			Array of preload lists (combined with items, crawl, and abr)
	items		Array of single items
	crawl		Array of directories to crawl
		prefix	Directory path
		depth	The level of depth to crawl
		acceptRegex	(Optional) Only download items with paths that match the regex.
		rejectRegex	(Optional) Do not download items with paths that match the regex.
	abr		Array of ABRs to download.
		manifest	Path of the manifest item. Only supports HLS, HSS, and DASH.
	retry		(Optional) Retry if download item fails. Default: off
		count	Max number of retries for a downloaded item that fails. Default: 2

Key1	Key2	Key3	Description
		delay	Number of seconds to delay before each retry. Default: 5 seconds
	schedule		(Optional) Scheduled time to download
		start	Time of day to start download. Should be the UTC time.
		interval	Amount of time to update the items. Mixed units with: "m" (minutes), "h" (hours) "d" (days). For example, "12d10h".

The following is an example of a simple Ingest Manifest file:

```
{
  "headers": "X-Auth-Key: mysecretkey",
  "preload" : [
    {
      "items": [
        "/content1/someone.jpg",
        "/content1/animals.mov"
      ],
      "retry": {
        "count": 2,
        "delay": 5
      },
      "schedule": {
        "start": "2016-06-25 08:25:25",
        "interval": "5d12h"
      }
    },
  ]
}
```

The following is an example of an ingest manifest file that references multiple single items, a directory crawl, and ABR:

```
{
  "headers": "X-Auth-Key: mysecretkey",
  "preload" : [
    {
      "items": [
        "/content1/someone.jpg",
        "/content1/animals.mov"
      ],
      "retry": {
        "count": 2,
        "delay": 5
      },
      "schedule": {
        "start": "2016-06-25 08:25:25",
        "interval": "5d12h"
      }
    },
    {
      "crawl": [
        {
          "prefix": "/wildfire/",
          "depth": 5,
          "rejectRegex": "\\..pl"
        },
        {

```

```
        "prefix": "/football/",
        "depth": 5,
        "acceptRegex": "\\\\.mov"
      }
    ],
    "abr": [
      {
        "manifest": "/nature/butterfly.m3u8",
      },
      {
        "manifest": "/nature/butterfly/Manifest",
      }
    ],
    "schedule": {
      "start": "2016-06-26 08:30:10",
      "interval": "2h30m"
    }
  }
}
```




APPENDIX C

NGB Whitelist File

This appendix describes the National Geo Blocking (NGB) whitelist file and provides an example of this file.

The NGB whitelist file is a JSON file that identifies a list of CIDR addresses that should be allowed even if NGB would block them. The NGB whitelist is assigned to a Traffic Router profile. Because only one Traffic Router profile can be applied to the Traffic Routers in a CDN, the NGB whitelist acts as a global list that will apply to all of the Delivery Services of the CDN.

The "NGB Whitelist File Format" table provides information on the format of the ingest manifest file. An example NGB whitelist file is shown after the table.

Table 6: NGB Whitelist File Format

Key	Description
date	Enter the date the file was created. This field is required but is currently not used.
name	Enter a descriptive name for the NGB whitelist file. This field is required but is currently not used.
version	This string can be used to track the version of the file. This field is required but is currently not used.
customer	Enter a string that describes the customers associated with this NGB whitelist file. This field is required but is currently not used.
ipWhiteList	Enter a JSON array of comma separated CIDR addresses that should be exempt from NGB blocking. This array can include a mix of IPv4 and IPv6 addresses in valid CIDR notation.

The following is an example of an NGB whitelist file:

```
{
  "date": "2018-08-10 22:59:03",
  "name": "National Geoblocking Whitelist Policy",
  "version": "1",
  "customer": "CU",
  "ipWhiteList": [
    "192.0.2.14/32",
    "192.0.2.158/32",
    "198.51.100.0/24",
    "203.0.113.0/24",
    "2001:DB8::1/64",
    "2001:DB8::2/64"
  ]
}
```

```
    ]  
}
```

For information on uploading the NGB whitelist file to the Traffic Router profile, see [NGB Whitelist, on page 176](#).



APPENDIX D

Header Rewrite Rules Syntax

This appendix describes the syntax for creating a header rewrite rule and provides some examples.

In addition to remapping a client request to an origin server, there may be a need to modify information in either the client request or the response header. For example, you may want to remove a particular header from a client request based on the header name and its value or you may want to change a customized error response that an Origin Server returns with a more generic error response.

The Header Rewrite settings of a Delivery Service enables you to define rules to modify the headers for both requests and responses. The rules that you define will determine which requests or responses will be modified, based on matching conditions, and how the header will be modified, based on the action (operator) you define.

Using OMD Director you can configure the following types of header rewrite rules:

- **Edge Header Rewrite Rules:** The Edge cache will modify either the request or response header, based on the rule you define.
- **Mid Header Rewrite Rules:** The Mid cache will modify either the request or response header, based on the rule you define.
- **Traffic Router Additional Response Headers:** The Traffic Router adds additional HTTP headers in the response to the client, based on the headers you define.

These header rewrite rules are configured on the Delivery Service and will apply to all of the Edge caches, Mid caches, or Traffic Routers that service that Delivery Service depending on the type of header rewrite rules you define.

A header rewrite rule consists of zero or more conditions followed by one or more operators. Conditions are used to limit the requests that will be affected by the operator(s):

- Conditions identify what criteria must be met before the operators will be evaluated.
- Operators determine what action will occur if the conditions are met.

Additionally, both conditions and operators can have flags that will modify their behavior.

This appendix has the following sections:

- [Conditions, on page 268](#)
- [Operators, on page 269](#)
- [Header Rewrite Rule Example, on page 269](#)
- [Configure Header Rewrite Rules, on page 269](#)

Conditions

The conditions part of the header rewrite rule have the following syntax:

cond %{<condition name> [: <argument>] } <operand> [<flags>]

- **cond**: Every condition begins with the literal string cond to indicate that this line is a condition, not an operator.
- **%{<condition name> [: <argument>] }**: cond is followed by the condition name, inside curly braces and preceded by a percent sign, for example %{HEADER} for the condition named HEADER. Some condition names take an argument. For example, header conditions take the name of the header to be matched and cookie conditions take the name of the cookie to be matched. For conditions that take arguments, the condition name is followed by a colon and then the argument value for example %{HEADER:Cache-Control}.
- **<operand>**: The operand provides a value, pattern, or range against the condition is matched. The following table describes the available operands:



Note

For some conditions the operand is optional. For these conditions if no operand is provided, any value that exists for that condition will cause the condition to be considered matched.

Table 7: Operands

Operand	Description
/regex/	Matches the provided value of the condition against the regular expression.
< "string"	Matches if the value from the condition is lexically less than the string value.
> "string"	Matches if the value from the condition is lexically greater than the string value.
= "string"	Matches if the value from the condition is lexically equal to the string value.

- **[<flags>]**: A condition may optionally have various flags associated with it, as described in the table below. Condition flags are optional and you can combine more than one condition into a comma-separated list of flags.

Table 8: Condition Flags

Flag	Description
AND	Indicates that both the current condition and the next must be true. This is the default behavior for all conditions when no flags are provided.
NOT	Inverts the condition.

Flag	Description
OR	Indicates that either the current condition or the next condition must be true.

Operators

Operators are the part of header rewrite rule that define how to modify the header content of the request or response. They are always the final part of a rule, following any of the conditions. You can specify multiple operators for a single rule. If more than one operator is specified, the operators are executed in the order listed.

For a complete list of conditions, operators, and their flags, please see https://docs.trafficserver.apache.org/en/6.2.x/admin-guide/plugins/header_rewrite.en.html.

Header Rewrite Rule Example

The following is an example of a rule that removes the Set-Cookie response header from a response:

```
cond %{READ_RESPONSE_HDR_HOOK}
rm-header Set-Cookie
```

Configure Header Rewrite Rules

You configure the header rewrite rules in the Advanced Settings window of the Delivery Service. For more information, see [Add a New DNS or HTTP Delivery Service, on page 92](#).



APPENDIX E

Configuring Enhanced DNS Request Routing

This appendix describes the Enhanced DNS Request Routing feature and how to configure it in Media Streamer Core Traffic Ops.

For Delivery Services that use DNS routing, clients do not make DNS requests directly to the Traffic Router. Typically the requests come from the DNS resolvers within the infrastructure. By default, the Traffic Router will use the IP address of the DNS resolver making the request to select the cache group to use for the request and not the client IP address.

RFC7871 defines an edns-client-subnet (ECS) EDNS0 option within the DNS query that contains the client subnet. You can configure the Traffic Router to use the client subnet in the ECS OPT field instead of the DNS resolver IP address to select a cache group. To configure the Traffic Router to use the client subnet in the ECS OPT field instead of the DNS resolver IP address, you must enable and configure the Enhanced DNS Routing Request (ECS) feature in the Traffic Router profile.

This appendix contains the following sections:

- [Configure Enhanced DNS Request Routing, on page 271](#)

Configure Enhanced DNS Request Routing

Perform the following steps to configure and enable the Enhanced DNS Routing Request feature (ECS) on the Traffic Router profile in Traffic Ops:

Procedure

- Step 1** Log in to the Traffic Ops UI:
`https://<ip or FQDN of the traffic ops>`
- Step 2** Choose **Parameters > Select Profile**.

Health	Delivery Services	Servers	Parameters	Tools
Search:	Profile	Host Name	Global Profile	Edge Cache Group
ALL	ALL	ALL	All Cache Groups	
EDGE_CDE250_CDN3	edge-cache2		All Profiles	test
EDGE_CDE250_CDN3	edge-cache3		Select Profile	test
EDGE_CDE250_CDN3	edge-cache4		Orphaned Parameters	test
ALL	ALL		EdgeCGTest	
ALL	ALL		MidCgTest	
Showing 1 to 6 of 6 entries				

Step 3 For the Traffic Router profile that you want to add the Enhanced DNS Routing Request feature (ECS) feature to, click **Parameter Details**.

Step 4 A list of parameters for the profile appears. To add a new parameter, click **Add Parameter**.

Health	Delivery Services	Servers	Parameters	Tools
Search:	Profile name	Parameter name		
CCR_CDN	CCR_CDN	anonymouisp.policy.configuration		
CCR_CDN	CCR_CDN	anonymouisp.polling.interval		
CCR_CDN	CCR_CDN	anonymouisp.polling.url		
CCR_CDN	CCR_CDN	api.cache-control.max-age		
CCR_CDN	CCR_CDN	api.port		
CCR_CDN	CCR_CDN	consistent.dns.routing		
CCR_CDN	CCR_CDN	coveragezone.polling.interval		
CCR_CDN	CCR_CDN	coveragezone.polling.url		
CCR_CDN	CCR_CDN	dnsec.dynamic.response.expiration		
CCR_CDN	CCR_CDN	domain.name		
CCR_CDN	CCR_CDN	federationmapping.polling.interval		
CCR_CDN	CCR_CDN	federationmapping.polling.url		
CCR_CDN	CCR_CDN	geolocation.polling.interval		
CCR_CDN	CCR_CDN	geolocation.polling.url		
CCR_CDN	CCR_CDN	keystore.maintenance.interval		
CCR_CDN	CCR_CDN	location		
CCR_CDN	CCR_CDN	monitor:///opt/tomcat/logs/access.log		
CCR_CDN	CCR_CDN	neustar.polling.interval		
CCR_CDN	CCR_CDN	neustar.polling.url		
CCR_CDN	CCR_CDN	proximity.cache.entries		
CCR_CDN	CCR_CDN	proximity.cache.timeout		
CCR_CDN	CCR_CDN	proximity.enabled		
CCR_CDN	CCR_CDN	proximity.ranking.depth		
CCR_CDN	CCR_CDN	proximity.server.dead.retry		
CCR_CDN	CCR_CDN	proximity.server.distributed_among_closest		
CCR_CDN	CCR_CDN	proximity.server.list		
CCR_CDN	CCR_CDN	ttl.soa.admin		
CCR_CDN	CCR_CDN	ttl.soa.expire		
CCR_CDN	CCR_CDN	ttl.soa.minimum		
CCR_CDN	CCR_CDN	ttl.soa.refresh		
CCR_CDN	CCR_CDN	ttl.soa.retry		
CCR_CDN	CCR_CDN	ttl.ttls.A		
CCR_CDN	CCR_CDN	ttl.ttls.AAAA		
CCR_CDN	CCR_CDN	ttl.ttls.DNSKEY		
CCR_CDN	CCR_CDN	ttl.ttls.DS		
CCR_CDN	CCR_CDN	ttl.ttls.HS		
CCR_CDN	CCR_CDN	ttl.ttls.SOA		
CCR_CDN	CCR_CDN	zonemanager.cache.maintenance.interval		
CCR_CDN	CCR_CDN	zonemanager.threadpool.scale		
Showing 1 to 39 of 39 entries				
Add Parameter				

Step 5 The Enter Parameter Details window appears. In this window, enter the following information:

- **Name:** ecsEnable
- **Config File:** CRConfig.json
- **Value:** True
- **Secure:** unchecked

Enter Parameter details

Basic Info:

* Name:

* Config File:

* Value:

* Secure: ☐

Step 6 Click **Save** to save the parameter.

Step 7 If you want to add the parameter to another Traffic Router profile, choose **Parameter > All Profiles**.

Step 8 In the **Search** field, enter **ecsEnable**. The list of profiles that currently use the ecsEnable parameter appears. Click **Edit** next to one of these profiles. The Parameter Detail window appears.

Step 9 In the Parameter Detail window, click **Add Profile** to add this parameter to another profile.

Parameter Detail

* Name:

* Config File:

* Value:

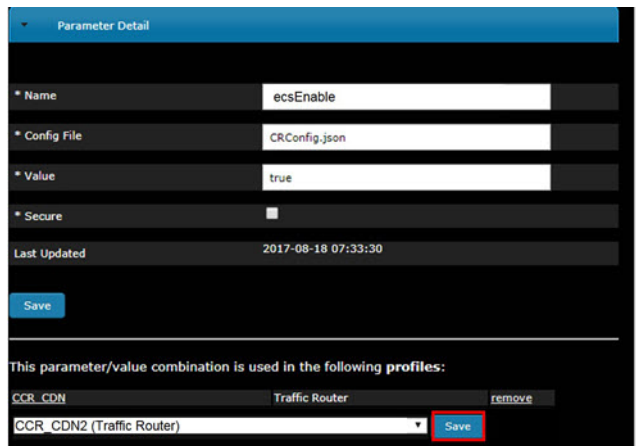
* Secure: ☐

Last Updated: 2017-08-18 07:33:30

This parameter/value combination is used in the following profiles:

CCR_CDN	Traffic Router	<input type="button" value="remove"/>
---------	----------------	---------------------------------------

Step 10 In the drop-down list that appears, choose the profile you want to add the ecsEnable parameter to and click **Save**.



The screenshot shows a 'Parameter Detail' form with the following fields and values:

Field	Value
Name	ecsEnable
Config File	CRConfig.json
Value	true
Secure	<input type="checkbox"/>
Last Updated	2017-08-18 07:33:30

Below the form is a 'Save' button. Underneath, a message states: 'This parameter/value combination is used in the following profiles:'. Below this message is a table with two columns: 'Profile' and 'Action'.

Profile	Action
CCR_CDN	Traffic Router
CCR_CDN2 (Traffic Router)	remove

A 'Save' button is located at the bottom right of the table, highlighted with a red box.

Step 11 Repeat Step 10 to add the ecsEnable parameter to additional Traffic Router profiles.



APPENDIX F

Recreate OMD Insights Summary Index Buckets

OMD Insights is based on Splunk, which uses buckets to store the indexed data that is used to create the charts and tables provided by OMD Insights. If you notice a gap in data in any of the charts or tables, or if you notice a sudden spike in any of the values, this indicates that a summary bucket may be missing or duplicated. This appendix describes how to delete existing summary index buckets and recreate them to resolve this issue.

To delete and recreate summary index buckets for a specific time range, perform the following steps:

1. Determine the time range for the summary index buckets that need to be deleted and recreated. You can determine this time range by determining where there are gaps or sudden spikes of data in the charts or tables.
2. Once you have determined the time range for the buckets that need to be deleted and recreated, log in to the OMD Insights Summary Search Heads as `ocdn_adm` and switch to root:

```
[ocdn_adm@omd-insights-sh ~]$ sudo su - root
```

3. Execute the following command, where `<start time>` and `<end time>` define the time range for the summary index buckets you need to delete and recreate, as determined in Step 1:

```
[root@omd-insights-sh ~]# /opt/splunk/bin/splunk cmd python Recreate_Summary_Buckets.py  
"<start time>" "<end time>"
```



Note

When executing this command it will prompt for the Splunk admin password, which is located in the `omd.conf` file.

For example, the following command would delete and recreate the summary index buckets for the time range from January 1 2018 to January 31 2018.

```
[root@omd-insights-sh ~]# /opt/splunk/bin/splunk cmd python Recreate_Summary_Buckets.py  
"01/01/2018:12:00:00" "02/01/2018:12:00:00"
```

4. Log out of the OMD Insights Summary Search Heads.
5. Confirm that the summary index buckets have successfully been deleted and recreated by viewing the chart or table during the time frame that originally had gaps or sudden spikes of data.



APPENDIX

G

tacreport tool

The tacreport tool collects detailed information about the hardware, software, and configuration of your Media Streamer system, and bundles it into a report file to transfer to Cisco support. All of the information that is gathered is considered confidential and Cisco will use this information only for diagnostic purposes. No changes will be made to your system during this process.

1. SSH into the system you are trying to analyze as ocdn_admin.

2. Change to the root user:

```
sudo su - root
```

3. Run the tacreport command.

```
tacreport
```



Note

This command normally takes a few minutes to complete, but may take longer if there are large log files. You can usually ignore any data collection error messages.

4. When the command is finished running, you will have a report file that you can provide to your Cisco account team.



APPENDIX H

Access and Transaction Log Details

This appendix describes the format and information found in the Traffic Router and Media Streamer cache server transaction logs. It also describes the cache response codes that may be in the Squid and Netscape log files.

This appendix includes the following sections:

- [Traffic Router Log Information, on page 279](#)
- [Media Streamer Cache Server Transaction Log Information, on page 285](#)

Traffic Router Log Information

The "Traffic Router Transaction Log Fields" table describes the fields that may be included in the Traffic Router `/opt/traffic_router/var/log/access.log` file.

Table 9: Traffic Router Transaction Log Fields

Name	Description	Data	Present
(#)	Epoch timestamp	The epoch time stamp of the request being logged	Always
qtype	Whether the request was for DNS or HTTP	Always DNS or HTTP	Always
chi	The IP address of the requester	Depends on whether this was a DNS or HTTP request	Always
ttms	The amount of time, in milliseconds, that it took the Traffic Router to process the request	A number greater than or equal to zero	Always

Name	Description	Data	Present
rtype	Routing result type		Always

Name	Description	Data	Present
		<p>This field will be one of the following:</p> <ul style="list-style-type: none"> • ANON_BLOCK: The request is being redirected (302) to the Anonymous Geo blocking URL. This blocking is based on the VPN/Hosting/Tor type IP addresses from a MaxMind database lookup, enabled on the Delivery service and defined by the <code>anonymousip.polling.url</code> setting for the Traffic Router profile. • CZ: The result was derived from Coverage Zone data, based on the address in the <code>chi</code> field. • DEEP_CZ: The result was derived from Deep Coverage Zone data, based on the address in the <code>chi</code> field. • DS_MISS_*HTTP Only*: No HTTP Delivery Service supports either the URL path or the headers of the request. • DS_REDIRECT: The result is using the Bypass Destination (Bypass FQDN) configured for the matched Delivery Service when that Delivery Service is unavailable or does not have the requested resource. • ERROR: An internal error occurred within the Traffic Router. More details may be found in the <code>rerr</code> field. • FED_*DNS Only*: The result was derived from federated coverage zone data outside of any Delivery Service. • GEO: The result was derived from the geolocation service, based on the address in the <code>chi</code> field. • GEO_REDIRECT: The request was redirected (302) based on the National Geo blocking (Geo Limit Redirect URL) configured on the Delivery Service. • MISS: The Traffic Router was unable to resolve a DNS request or find a cache for the requested resource. • RGALT: The request was redirected (302) to the Regional Geo blocking URL. Regional Geo blocking is enabled on the Delivery Service and is configured through the <code>regional_geoblock.polling.url</code> setting for the 	

Name	Description	Data	Present
		<p>Traffic Router profile.</p> <ul style="list-style-type: none">• RGDENY: The request was regionally blocked because there was no rule for the request made.• STATIC_ROUTE_*DNS Only*_: No DNS Delivery Service supports the hostname portion of the requested URL.• -: The request was not redirected. This is usually a result of a DNS request to the Traffic Router or an explicit denial for that request.	
rloc	Geolocation of result	Latitude and longitude in decimal degrees. (This will be “-” for DNS.)	Always

Name	Description	Data	Present
rdtl	Result details associated with unusual conditions	<p>This field will be one of the following:</p> <ul style="list-style-type: none"> • DS_NOT_FOUND: Always goes with the rtypes of STATIC_ROUTE and DS_MISS. • DS_BYPASS: The Bypass Destination (Bypass FQDN) configured for the matched Delivery Service was used to redirect the request. • DS_CLIENT_GEO_UNSUPPORTED: The Traffic Router did not find a resource supported by the Coverage Zone data and was unable to determine the geolocation of the requesting client. • DS_CZ_ONLY: The selected Delivery Service only supports resource lookup based on Coverage Zone data. • DS_NO_BYPASS: No valid Bypass Destination (Bypass FQDN) is configured for the matched Delivery Service and the Delivery Service does not have the requested resource. • GEO_NO_CACHE_FOUND: The Traffic Router could not find a resource using the geolocation data based on the geolocation of the requesting client. • NO_DETAILS: There is no unusual request being made. This entry is for a standard request. • REGIONAL_GEO_ALTERNATE_WITHOUT_CACHE: This goes with the rtype RGDENY. The URL is being regionally Geo blocked. • REGIONAL_GEO_NO_RULE: The request was blocked because there was no rule in the Delivery Service for the request. • -: The request was not redirected. This is usually a result of a DNS request to the Traffic Router or an explicit denial for that request. 	Always
rerr	Message about an internal Traffic Router error	String (if error exists)	Always
url	Requested URL with query string	String	HTTP
cqhm	HTTP method in request	HTTP method in request, such as GET or POST	HTTP
cqhv	HTTP version in request	For example: HTTP/1.1	HTTP

Name	Description	Data	Present
rh	Key value pair. One or more key value pairs may exist in a logged event and are controlled by the configuration of the matched Delivery Service.	Key value pair of the format “name: value”	HTTP
rurl	The resulting URL of the resource requested by the client	A URL string	HTTP
rgb	Postal code and RGB lookup information	If a URL has been defined for Regional Geo Blocking, the details provided in this field will reflect the postal code and additional information.	HTTP
pssc	Traffic router response code	How the traffic router responded to the HTTP request, for example 302, 503, etc. For a complete list of possible codes refer to https://en.wikipedia.org/wiki/List_of_HTTP_status_codes	HTTP
xn	The ID from the client DNS request header	A number from 0 to 65535	DNS
fqdn	The qname field from the client DNS request message, for example, the FQDN the client is requesting to be resolved	A series of DNS labels or domains separated by ‘.’ characters and ending with a ‘.’ character	DNS
type	The qtype field from the client DNS request message, for example, the type of resolution that is requested such as IPv4, IPv6	Examples for this field are A (IPv4), AAAA (IPv6), NS (Name Service), SOA (Start of Authority), and CNAME	DNS
class	The qclass field from the client DNS request message, for example, the class of resource being requested	This field can be either IN (Internet resource) or ANY (Traffic router rejects requests with any other value of class).	DNS
ttl	The Time to Live, in seconds, for the answer from the Traffic Router. This is how long clients can reliably use the reply without re-querying the Traffic Router.	A number from 0 to 4294967295	DNS
rcode	The result code for the DNS answer that was provided by the Traffic Router	This field will be one of the following: <ul style="list-style-type: none"> • NOERROR: Success • NOTIMP: Request is not supported • REFUSED: Request is refused to be answered • NXDOMAIN: The domain/name requested does not exist 	DNS

Name	Description	Data	Present
ans	IP address in answer to the query performed	IP address in answer to the query performed	DNS

For additional information, refer to http://traffic-control-cdn.readthedocs.io/en/latest/admin/traffic_router.html#event-log-file-format.

Media Streamer Cache Server Transaction Log Information

The "Media Streamer Cache Server Transaction Log Fields" table describes the fields that may be included in the Media Streamer cache servers transaction log files. The transaction log file is named openmd_transaction.log and is located in the /opt/trafficserver/var/log/trafficserver/ directory on each cache server.

Table 10: Media Streamer Cache Server Transaction Log Fields

Name	Description	Data
chi	The IP address of the client host machine	IP address
phn	The hostname of the Traffic Server that generated the log entry in collated log files	Hostname
url	URL being requested	A URL string
cqhm	The HTTP method in the client request to the Traffic Server. Note Not all HTTP request methods will be honored in the CDN.	For example: GET or HEAD
cqhv	The HTTP version of the client request	For example: HTTP/1.1
pssc	The HTTP response status code from the Traffic Server to the client	Any HTTP response that is applicable to the operation, for example, 200, 302, 404, etc. For a complete list of possible codes refer to https://en.wikipedia.org/wiki/List_of_HTTP_status_codes
ttms	The time the Traffic Server spends processing the client request. This is calculated as the number of milliseconds between the time the client establishes the connection with the Traffic Server and the time the Traffic Server sends the last byte of the response back to the client.	The value in milliseconds
b	Response length of the content	The value in bytes
sssc	The HTTP response status code from the origin server to the Traffic Server	Any HTTP response that is applicable to the operation, for example 200, 302, 404, etc. If the content is cached and has not timed out this will be 0000.

Name	Description	Data
sscl	The response length (in bytes) from the origin server to the Traffic Server	The value in bytes. If the content is cached and has not timed out this will be 0.
cfsc	The client finish status code	Will be one of the following: <ul style="list-style-type: none"> • FIN: The client request to the Traffic Server was successfully completed. • INTR: The client request to the Traffic Server was interrupted.
pfsc	The proxy finish status code	Will be one of the following: <ul style="list-style-type: none"> • FIN: The Traffic Server request to the origin server was successfully completed. • INTR: The Traffic Server request to the origin server was interrupted. • TIMEOUT: The Traffic Server request to the origin server timed out. Mid cache servers are proxies on behalf of the Edge cache nodes.
crc	The cache result code	Specifies how the cache responded to the request, for example HIT, MISS, etc. For a full list of possible responses, refer to https://docs.trafficserver.apache.org/en/6.2.x/admin-guide/monitoring/logging/log-formats.en.html#admin-monitoring-logging-cache-result-codes
phr	The proxy hierarchy route, which is the route the Traffic Server used to retrieve the object	<ul style="list-style-type: none"> • PARENT_HIT • NONE: For cache hit • Direct: Response from local system (astats) or single origin (no MSO)
phi	The primary IP address of the Traffic Server that generated the log entry in collated log files	IP address
hii	The IP address the Traffic Server used to stream the content. This is the IP address to which the client connected.	IP address
pqsn	The IP address or FQDN of the proxy request server	<ul style="list-style-type: none"> • 0 on cache hits • The parent-ip for requests to parent proxies.
uas	User agent string	Client, client OS, or platform specifics

Name	Description	Data
psqn	Origin server configured FQDN	For edge caches this will be "-". For Mid caches, this will be the FQDN of the origin server that is configured for the delivery service from which the content was requested. If MSO is configured, this will be one of the configured MSO origin servers for the delivery service.
shn	The hostname of the origin server for the delivery service.	This is the FQDN for the origin server configured for the delivery service. Note, if MSO is configured and the access.log is from a MID cache node, the shn is the origin server defined as the lookup for the delivery service. The actual origin server from which the content is requested is the psqn.

For additional information refer to:

- For Apache Traffic Server 5.3.2: <https://docs.trafficserver.apache.org/en/5.3.x/admin/event-logging-formats.en.html>
- For Apache Traffic Server 6.2: <https://docs.trafficserver.apache.org/en/6.2.x/admin-guide/monitoring/logging/log-formats.en.html>



APPENDIX

OMD Director Alarms and Remediation

This Appendix describes the default checks and their thresholds that are run by the sensu-client service on the OMD Monitor clients that you can view and deactivate using OMD Director. To view these checks and their thresholds from OMD Director, choose **Monitor > Alarms** and click the **Alarm Rules** tab. This Appendix also describes possible remediation steps that you can perform if an alarm is raised.

When you triage the alerts that are raised, you should check the following items:

- Any active network configuration change requests that could affect traffic
- The active Media Streamer change requests for the servers in question
- Whether the alert indicates that end customer services are impacted, such as the alert is on a Cache node or the OMD Insights server is down
- Run the **dmesg | grep -i error** command and check for any errors
- Check the `/var/log/messages` files for any errors (`grep -i error`)
- Check OMD Insights to determine whether the effected node provided content during the error window

All notifications are generated based on the thresholds that are configured for the alarms. An alarm can have multiple thresholds configured for it to send notifications about different severity levels. To view the existing thresholds and to create new thresholds, from the **Monitor > Alarms** page, click the **Alarm Rules** tab. The "Checks and Thresholds Available in OMD Director" table describes the default checks and thresholds that are available in OMD Director. [Table 12: Alarm Remediation Steps , on page 299](#) describes possible remediation steps you can perform based on the alarms or alerts that are raised.

Table 11: Checks and Thresholds Available in OMD Director

Check Name	Entity Family (entity_family in Influxdb)	Check Value	Alarm Condition/ Check Result [0 - No Alert/Normal] [1 - Warning Alert] [2 - Critical Alert]	Severity/Priority
bond_interface	cache_nodes	Status of the bond interface: <ul style="list-style-type: none">• 0: If the bond interface is up and running• 1: If the bond interface is not configured• 2: If the bond interface is down	Warning: Alert raised if bond interface is not configured Critical: Alert raised if bond interface is configured and down	P4
cpu_usage	common	CPU usage in percentage. (0 - 100%)	Warning alarm raised if: <ul style="list-style-type: none">• $\geq 50\%$ (Default)• $\geq 80\%$ (Monitor Node)• $\geq 55\%$ (Edge) Critical alarm raised if: <ul style="list-style-type: none">• $\geq 80\%$ (Default)• $\geq 95\%$ (Monitor Node)• $\geq 85\%$ (Edge)	Warning: N/A Critical: P4
disk_drives_count	cache_nodes	Count of current disk partitions	Critical: Alert raised if the number of disk drive partitions currently available is less than the original number of partitions	P3

Check Name	Entity Family (entity_family in Influxdb)	Check Value	Alarm Condition/ Check Result [0 - No Alert/Normal] [1 - Warning Alert] [2 - Critical Alert]	Severity/Priority
disk_usage	common	Total disk usage in percentage	Warning alarm raised if: <ul style="list-style-type: none"> • >= 85% (Default) • >= 90% (Monitor Node) Critical alarm raised if: <ul style="list-style-type: none"> • >= 95% 	Warning: N/A Critical: P3/P4
dns_resolved_status	common	Status of DNS resolution: <ul style="list-style-type: none"> • 0: DNS is not resolvable • 1: DNS is resolvable 	Critical: Alert raised if DNS is not resolvable (check result returns 0)	P4
hardware	common	Hardware error message status from the dmesg command output: <ul style="list-style-type: none"> • 0: No errors found • 1: Errors found 	Critical: Hardware error reported in dmesg.	P3
influx_connectivity	common	Checks connectivity to OMD Monitor Influx nodes. <ul style="list-style-type: none"> • 0: Normal connectivity • 2: No connectivity to any OMD Monitor Influx nodes 	Critical: Alert raised if none of the OMD Monitor Influx nodes are reachable.	P3
influx_metrics_not_found	common	<ul style="list-style-type: none"> • 0: Normal no errors found • 2: Critical 	Critical: Metrics not found from OMD Monitor Influx nodes at this time	

Check Name	Entity Family (entity_family in Influxdb)	Check Value	Alarm Condition/ Check Result [0 - No Alert/Normal] [1 - Warning Alert] [2 - Critical Alert]	Severity/Priority
insights_failover_check	splunk	<ul style="list-style-type: none"> • 0: Failover did not happen and the primary OMD Insights Splunk nodes are active. • 1: Failover happened for the Splunk service on the backup OMD Insights Splunk nodes 	Warning: Failover occurred, the standby node is currently the master.	P4
interface	cache_nodes	Status of the slave interfaces of the bond interface: <ul style="list-style-type: none"> • 0: If the interface is up and running • 1: If the interface is not configured • 2: If the interface is down 	Warning: Alert raised if an interface that is part of the bond interface is not configured Critical: Alert raised if an interface that is part of the bond interface is down	P3
kafka_broker_status_all	monitor-node	The number of servers that are having a problem with the Kafka broker server. 0 means all servers are fine or there are no servers configured for Kafka export.	Critical: Alert raised if any or all Kafka brokers are not able to export metrics	P3/P4
keep_alive	common	The status of the Monitor node connection to the server node being monitored. To determine this, keepalives are sent by the sensu-client service on the server node to a Monitor node every 30 seconds: <ul style="list-style-type: none"> • 0: Not reachable • 1: Reachable 	Warning: Alert raised if no keepalives have been received for 90 to 180 seconds Critical: Alert raised if no keepalives have been received for more than 180 seconds	P3/P2 The priority is determined based on the number of servers that are raising the keepalive. If a single server is in this state, it is a P3. If there is a block of servers that are affected, it is a P2 and you should escalate immediately.

Check Name	Entity Family (entity_family in Influxdb)	Check Value	Alarm Condition/ Check Result [0 - No Alert/Normal] [1 - Warning Alert] [2 - Critical Alert]	Severity/Priority
load_avg_1min	common	CPU load average values, with regards to both the CPU and I/O over the last minute. This value is a decimal number and not a percentage. These values are taken from the /proc/loadavg system file.	Warning: Alert raised if load_avg_1min >= 10 Critical: Alert raised if load_avg_1min >= 25	Warning: N/A Critical: P4
load_avg_5min	common	CPU load average values, with regards to both the CPU and I/O over the last 5 minutes. This value is a decimal number and not a percentage. These values are taken from the /proc/loadavg system file.	Warning: Alert raised if load_avg_5min >= 20 Critical: Alert raised if load_avg_5min >= 50	Warning: N/A Critical: P4
load_avg_15min	common	CPU load average values, with regards to both the CPU and I/O over the last 15 minutes. This value is a decimal number and not a percentage. These values are taken from the /proc/loadavg system file.	Warning: Alert raised if load_avg_15min >= 30 Critical: Alert raised if load_avg_15min >= 75	Warning: N/A Critical: P4
log_failed_notallowed	common	The number of SSH failed login attempts due to access restrictions	Warning: Alert raised if 1 failed attempt Critical: Alert raised if 2 or more failed attempts	P4
log_failed_password	common	The number of failed login attempts do to incorrect passphrase	Warning: Alert raised if 1 failed attempt Critical: Alert raised if 2 or more failed attempts	P4

Check Name	Entity Family (entity_family in Influxdb)	Check Value	Alarm Condition/ Check Result [0 - No Alert/Normal] [1 - Warning Alert] [2 - Critical Alert]	Severity/Priority
netstat	ESTABLISHED	<p>This check indicates the number of connections in the ESTABLISHED state.</p> <p>An alarm is raised if the number of connections exceeds the configured threshold. A server only has a finite number of TCP connections available, so this alarm indicates when the server is nearing that threshold.</p> <p>Connections are determined by looking for ESTABLISHED and TIME_WAIT entries using the netstat command.</p>	<p>Warning: Alert is raised if ESTABLISHED connections are at the following thresholds:</p> <ul style="list-style-type: none"> • >= 500 (Default) • >= 40000 (cache nodes) • >= 700 (tcnode) • >= 1500 (monitor node) <p>Critical: Alert is raised if ESTABLISHED connections are at the following thresholds:</p> <ul style="list-style-type: none"> • >= 1000 (Default) • >= 45000 (cache nodes) • >= 1200 (tcnode) • >= 2000 (monitor node) 	<p>Warning: N/A</p> <p>Critical: P4</p>

Check Name	Entity Family (entity_family in Influxdb)	Check Value	Alarm Condition/ Check Result [0 - No Alert/Normal] [1 - Warning Alert] [2 - Critical Alert]	Severity/Priority
netstat	TIME_WAIT	<p>This check indicates the number of connections in the TIME_WAIT state.</p> <p>An alarm is raised if the number of connections exceeds the configured threshold. A server only has a finite number of TCP connections available, so this alarm indicates when the server is nearing that threshold.</p> <p>Connections are determined by looking for ESTABLISHED and TIME_WAIT entries using the netstat command.</p>	<p>Warning: Alert is raised if ESTABLISHED connections are at the following thresholds:</p> <ul style="list-style-type: none"> • >= 500 (Default) • >= 40000 (cache nodes) • >= 700 (tcnode) • >= 1500 (monitor node) <p>Critical: Alert is raised if ESTABLISHED connections are at the following thresholds:</p> <ul style="list-style-type: none"> • >= 1000 (Default) • >= 45000 (cache nodes) • >= 1200 (tcnode) • >= 2000 (monitor node) 	<p>Warning: N/A</p> <p>Critical: P4</p>
ntp_sync_status	common	The status of the sync with the NTP server. 0 indicates the server is not in sync. 1 indicates the server is in sync.	Critical alarm raised if ntp_offset >= 500ms	Critical: P4

Check Name	Entity Family (entity_family in Influxdb)	Check Value	Alarm Condition/ Check Result [0 - No Alert/Normal] [1 - Warning Alert] [2 - Critical Alert]	Severity/Priority
partition_usage	Name of each partition: /var /sys/fs/cgroup /run /opt /dev /dev/shm /boot /	Disk usage for each partition, as a percentage. This check will return 1(warning) if the configured warning or critical threshold values are more than 100.	Warning: Alert raised if any partition usage >= 85% Critical: Alert raised if any partition usage >= 95%	Warning: N/A Critical: P4
ram_usage	common	RAM usage as a percentage	Warning: Alert raised if RAM Usage is >= 80% Critical: Alert raised if RAM Usage is >= 95%	P4
redis_status	monitor-node	The number of connected OMD Monitor nodes. This number should be equal to the number of OMD Monitor nodes available at install. If it is not, a warning is raised.	Warning: Alert raised if the number of connected OMD Monitor nodes is less than the number of OMD Monitor nodes available at initial installation.	N/A
service_status	grafana-server	The running status of the Grafana server on the OMD Monitor node: • 0: Not running • 1: Running	Critical: Alert raised when the grafana-server service does not have the status of “running”.	P4
service_status	haproxy	The running status of the haproxy service on the OMD Monitor node: • 0: Not running • 1: Running	Critical: Alert raised when the haproxy service does not have the status of “running”.	P4

Check Name	Entity Family (entity_family in Influxdb)	Check Value	Alarm Condition/ Check Result [0 - No Alert/Normal] [1 - Warning Alert] [2 - Critical Alert]	Severity/Priority
service_status	influxdb	The running status of the influxdb service on the OMD Monitor node: <ul style="list-style-type: none">• 0: Not running• 1: Running	Critical: Alert raised when the influxdb service does not have the status of “running”.	P4
service_status	influxdb-relay	The running status of the influxdb-relay service on the OMD Monitor node: <ul style="list-style-type: none">• 0: Not running• 1: Running	Critical: Alert raised when the influxdb-relay service does not have the status of “running”.	P4
service_status	rabbitmq-metrics-exporter	Running status of rabbitmq-metrics-exporter on monitor node <ul style="list-style-type: none">• 0: Running• -1: Not Running	Critical when rabbitmq-metrics-exporter service is not in a running state.	P2/P3
service_status	rabbitmq-server	The running status of the rabbitmq-server service on the OMD Monitor node: <ul style="list-style-type: none">• 0: Not running• 1: Running	Critical: Alert raised when the rabbitmq-server service does not have the status of “running”.	P4/P3
service_status	redis	The running status of redis service on the OMD Monitor node: <ul style="list-style-type: none">• 0: Not running• 1: Running	Critical: Alert raised when the redis service does not have the status of “running”.	P4/P3
service_status	redis-sentinel	The running status of redis-sentinel service on the OMD Monitor node: <ul style="list-style-type: none">• 0: Not running• 1: Running	Critical: Alert raised when the redis-sentinel service does not have the status of “running”.	P4/P3

Check Name	Entity Family (entity_family in Influxdb)	Check Value	Alarm Condition/ Check Result [0 - No Alert/Normal] [1 - Warning Alert] [2 - Critical Alert]	Severity/Priority
service_status	sensu-server	The running status of sensu-server service on the OMD Monitor node: <ul style="list-style-type: none">• 0: Not running• 1: Running	Critical: Alert raised when the sensu-server service does not have the status of “running”.	P4/P3
service_status	splunk	The running status of splunk service on the OMD Monitor node: <ul style="list-style-type: none">• 0: Not running• 1: Running	Critical: Alert raised when the splunk service does not have the status of “running”.	P4/P3
service_status	tcp-rabbitmq-exchange	Running status of tcp-rabbitmq-exchange on monitor node: <ul style="list-style-type: none">• 0: Running• -1: Not Running	Critical when tcp-rabbitmq-exchange service is not in a running state.	P2/P3
sshd_running_status	common	The running status of the sshd service: <ul style="list-style-type: none">• 0: Not running• 1: Running	Warning: Alert raised if the sshd service is not running.	P4
swap_usage	common	The percentage of swap usage	Warning: Alert raised when swap usage is >= 95% Critical: Alert raised when swap usage is >= 98%	P2

Check Name	Entity Family (entity_family in Influxdb)	Check Value	Alarm Condition/ Check Result [0 - No Alert/Normal] [1 - Warning Alert] [2 - Critical Alert]	Severity/Priority
tcp_connection	common cache_nodes traffic_router	Number of TCP connections including TCP6	Warning: Alert raised if TCP connections are: <ul style="list-style-type: none"> • >= 500 (others) • >= 50000 (Cache nodes) • >= 5000 (Traffic router) Critical: Alert raised if TCP connections are: <ul style="list-style-type: none"> • >= 1000 (others) • >= 55000 (Cache nodes) • >= 10000 (Traffic router) 	Warning: N/A Critical: P4
tripwire_violations	common	The number of tripwire report violations. Local local files are checked for changes.	Critical: If any tripwire violations are found (check result is 3)	P4

Table 12: Alarm Remediation Steps

Check	Remediation
bond_interface	If this alarm has been raised, the alarm should have already been cleared so a keepalive alarm should have been raised. Also look at the IPMI Interface check in Traffic Ops to see if the server has power.

Check	Remediation
cpu_usage	<p>Level 3:</p> <ul style="list-style-type: none"> • Check the running state of the system to determine if there is an issue. If the alert is generated and resolved immediately, it may be able to be ignored. • This may or may not be an issue depending on the load/traffic on the server. If the alert continuously persists, execute the top command on the client terminal to determine which process is running high on CPU. • Check the Grafana dashboard for any patterns in the usage.
disk_drives_count	<p>Log in to the cache node and use the ls /dev/sd* command to check the disk drive partitions list.</p> <p>Based on the profile of Mid or Edge cache sever, the list will change. sda, sdb, sdc, and so on are the disk drives and sda1, sda2, and so on are partitions in sda. Check if any disk is missing based on the profile.</p>
disk_usage	<p>Level 2/Level3: For critical alerts, log in and check how different partitions are being used. Use the df command to determine which partitions have high disk usage.</p>
dns_resolved_status	<p>Level 1/Level 2: Make sure servers have access to DNS servers and that those DNS servers are correct.</p>

Check	Remediation
hardware	<p>Indicates a possible hardware error. Details can be found in the <code>/var/log/messages</code> file. The check is looking for “Hardware Error” in the dmesg command output. Note that:</p> <ul style="list-style-type: none"> • If “hardware error” is seen in the dmesg output, but has not been corrected in the mcelog (located in <code>/var/log/messages</code>), an alert will be raised immediately. • Alerts are not raised if dmesg lists the same error based on the dmesg time stamp. • Alerts are not raised for corrected errors shown in the mcelog (located in <code>/var/log/messages</code>) • Alerts are raised if the correctable errors are reported more than 6 hours of the day. <p>Perform the following steps:</p> <ol style="list-style-type: none"> 1. If the alert is repeatedly raised, then Stash (mute) the alert while investigating the issue. 2. SSH to the affected system and change to the root user. 3. Collect the relevant portion of the <code>/var/log/messages</code> file for information on the hardware error. 4. Send an email for a case to be opened, including all of the collected information. See the note below. 5. Back up the dmesg log using the following command: dmesg > /var/log/dmesg-`date +%Y%m%d_%I%M%S` 6. Enter the dmesg -c command to clear the dmesg log. 7. Unstash (unmute) the alert. <p>Note Level 1 should create a case with the hardware vendor and ask them to identify what the error means and to recommend the next steps. If the vendor recommends a DIMM replacement, ensure that the issue has actually been identified. Ask the vendor to explain why the DIMM needs to be replaced and to identify which slot is referenced in the error message.</p>
influx_connectivity	<p>Check whether there is a network event that is causing the loss of connectivity. If no possible cause can be found, escalate to Level 2 or Level 3.</p>
insights_failover_check	<p>Check the connectivity between the primary and backup OMD Insights nodes. There may be network connectivity issues between the primary and backup OMD Insights nodes that resulted in this alarm.</p> <p>Check for any keepalive alerts from OMD Monitor for the OMD Insights primary nodes.</p>

Check	Remediation
interface	Level 2: Verify that the issues is not external to the CDN cache nodes and then escalate as needed.
kafka_broker_status_all	<p>Check the accessibility of the Kafka servers from the OMD Monitor nodes:</p> <ul style="list-style-type: none"> • Check the omd.conf pillar file on the Salt Master and verify the Kafka servers and their ports are configured correctly. • SSH to the OMD monitor node and verify connectivity to the Kafka servers. • Check the /var/log/sensu/sensu-server.log and /etc/sensu/kafka_broker_status.log log files for any error messages.
keep_alive	<p>Network connectivity between an OMD Monitor node and the server being monitored might have been lost. Additional insight into the network is required.</p> <p>If you can access the server using SSH, you can run the following commands to determine whether the system rebooted or if there was a network issue:</p> <ul style="list-style-type: none"> • uptime: This command shows how long the system has been up. If it has been up for well before the alert, than use the next command to further troubleshoot. • grep “NIC Link” /var/log/messages*: This will display the details for the interface status, if it had a problem. <p>If neither of these commands reveal an issue, perform additional review of the following log files:</p> <ul style="list-style-type: none"> • /var/log/sensu/sensu-client.log on the server • /var/log/sensu/sensu-server.log on Monitor nodes
load_avg_1min	<p>Level 3:</p> <ul style="list-style-type: none"> • Check the running state of the system to determine if there is an issue. If the alert is generated and resolved immediately, it may be able to be ignored. • This may or may not be an issue depending on the load/traffic on the server. If the alert continuously persists, use the top command from the client terminal to determine which process is running high on CPU. • Check the Grafana dashboard for any patterns in the usage.

Check	Remediation
load_avg_5min	<p>Level 3:</p> <ul style="list-style-type: none"> • Check the running state of the system to determine if there is an issue. If the alert is generated and resolved immediately, it may be able to be ignored. • This may or may not be an issue depending on the load/traffic on the server. If the alert continuously persists, use the top command from the client terminal to determine which process is running high on CPU. • Check the Grafana dashboard for any patterns in the usage.
load_avg_15min	<p>Level 3:</p> <ul style="list-style-type: none"> • Check the running state of the system to determine if there is an issue. If the alert is generated and resolved immediately, it may be able to be ignored. • This may or may not be an issue depending on the load/traffic on the server. If the alert continuously persists, use the top command from the client terminal to determine which process is running high on CPU. • Check the Grafana dashboard for any patterns in the usage.
log_failed_notallowed	<p>Level 3: SSH to the server as root and check the /var/log/secure file for a string similar to the following:</p> <pre>"User root from 11.22.33.44 not allowed because not listed in AllowUsers"</pre>
log_failed_password	<p>SSH to the server as root and check the /var/log/secure file for the string "Failed password".</p>
netstat for ESTABLISHED	<p>SSH to the server. Check the output from the netstat command to determine more information about the connections and which port is being used. Look for TCP connections in the ESTABLISHED states, which could provide information about the related service.</p>
netstat for TIME_WAIT	<p>SSH to the server. Check the output from the netstat command to determine more information about the connections and which port is being used. Look for TCP connections in the TIME_WAIT states, which could provide information about the related service.</p>
ntp_sync_status	<p>Level 2/Level 3: Connectivity to the NTP server needs to be checked or other network resources are not responding. Use the ntpq -p command to check the NTP source stratum for correct time.</p> <p>Level 3 should fix the NTP configuration as needed, or perform additional troubleshooting to ensure that NTP is synced.</p>
partition_usage	<p>Level 2/Level 3: SSH to the server or view Grafana dashboards for the server to determine which partition is high on usage. After connecting to the server using SSH, look for files that are consuming a large amount of disk space.</p>

Check	Remediation
ram_usage	<p>SSH to the server and check the memory usage. Use commands such as <code>top</code> to identify the processes that are running and determine which processes are consuming memory.</p> <p>Check Grafana dashboards for any patterns in the RAM usage.</p>
redis_status	<p>A warning for this check indicates that an OMD Monitor node could not connect to one or more of the other OMD Monitor node. However, if OMD Monitor is installed in an HA configuration, the OMD Monitor nodes should continue to work without issue.</p> <p>You should investigate the reason for the loss of connection between OMD Monitor nodes. If there is no issue with connectivity, review the <code>/var/log</code> files on the OMD Monitor node.</p>
service_status for grafana-server	<p>Connect to the OMD Monitor node and enter the <code>systemctl status grafana-server</code> command to check the status of the grafana-server service. If the service is not in a running state, enter the <code>sudo systemctl restart grafana-server</code> command to try to bring it up.</p> <p>Note If the grafana-server is not in a running state, the Grafana GUI will also fail to load when connected to this OMD Monitor node.</p> <p>Check the <code>/var/log/grafana/grafana.log</code> file for any error messages.</p> <p>If the alarm persists, escalate the problem for further investigation.</p>
service_status for haproxy	<p>Connect to the OMD Monitor node and enter the <code>systemctl status haproxy</code> command to check the status of the haproxy service. If the service is not in a running state, enter the <code>sudo systemctl restart haproxy</code> command to try to bring it up.</p> <p>Check the <code>/var/log/haproxy.log</code> and <code>/var/log/haproxy-status.log</code> files for any error messages.</p> <p>If the alarm persists, escalate the problem for further investigation.</p>
service_status for influxdb	<p>Connect to the OMD Monitor node and enter the <code>systemctl status influxdb</code> command to check the status of the influxdb service.</p> <p>Influxdb metrics may have stopped on one OMD Monitor node, however, the remaining OMD Monitor nodes will continue to store metrics.</p> <p>Enter the <code>sudo systemctl restart influxdb</code> command to try to bring up the influxdb service on OMD Monitor node. If the alarm persists, escalate the problem for investigation.</p>

Check	Remediation
service_status for influxdb-relay	<p>Connect to the OMD Monitor node and enter the systemctl status influxdb-relay command to check the status of the influxdb-relay service.</p> <p>Influxdb metrics may have stopped on one OMD Monitor node, however, the remaining OMD Monitor nodes will continue to store metrics.</p> <p>Enter the sudo systemctl restart influxdb-relay command to try to bring up the influxdb-relay service on the OMD Monitor node. If the alarm persists, escalate the problem for investigation.</p>
service_status for rabbitmq-metrics-exporter	<p>The rabbitmq-metrics-exporter service pops out metrics data from RabbitMQ and exports the same to the Monitor Nodes, InfluxDBs, and the external Kafka servers. Failure of this service would result in no metrics data in Influxdb but the Monitor mailer alerts and Uchiwa will still be functional.</p> <p>Connect to the Monitor node and check the status of the process using <code>systemctl status rabbitmq-metrics-exporter</code>.</p> <p>Inspect <code>/var/log/messages</code> on Monitor node for any error messages.</p> <p>Try bringing up the service using <code>sudo systemctl restart rabbitmq-metrics-exporter</code> on the Monitor node. If the alarm persists, then escalate for investigation.</p>
service_status for rabbitmq-server	<p>Connect to the OMD Monitor node and enter the systemctl status rabbitmq-server command to check the status of the rabbitmq-server service.</p> <p>An OMD Monitor node whose rabbitmq-server service is not running will not be monitoring any clients. Clients previously connected to this node will move to another OMD Monitor node, including the Sensu client of the OMD Monitor node whose rabbitmq-server service is not running.</p> <p>Inspect the <code>/var/log/rabbitmq/</code> logs for any error messages.</p> <p>Enter the sudo systemctl restart rabbitmq-server command to try to bring up the rabbitmq-server service on the OMD Monitor node. If the alarm persists, escalate the problem for investigation.</p>
service_status for redis	<p>Connect to the OMD Monitor node and enter the systemctl status redis command to check the status of the redis service.</p> <p>An OMD Monitor node whose redis service is not running will not be monitoring any clients. Clients previously connected to this node will move to another OMD Monitor node, including the Sensu client of the OMD Monitor node whose redis service is not running.</p> <p>Inspect the <code>/var/log/redis/</code> logs for any error messages.</p> <p>Enter the sudo systemctl restart redis command to try to bring up the redis service on the OMD Monitor node. If the alarm persists, escalate the problem for investigation.</p>

Check	Remediation
service_status for redis-sentinel	<p>Connect to the OMD Monitor node and enter the systemctl status redis-sentinel command to check the status of the redis-sentinel service.</p> <p>Inspect /var/log/redis/ logs for any error messages.</p> <p>Enter the sudo systemctl restart redis-sentinel command to try to bring up the redis-sentinel service on the OMD Monitor node. If the alarm persists, escalate the problem for investigation.</p>
service_status for sensu-server	<p>Connect to the OMD Monitor node and enter the systemctl status sensu-server command to check the status of the sensu-server service.</p> <p>Inspect the /var/log/sensu/sensu-server.log file for any error messages.</p> <p>Enter the sudo systemctl restart sensu-server command to try to bring up the sensu-server service on the OMD Monitor node. If the alarm persists, escalate the problem for investigation.</p>
service_status for splunk	<p>The Splunk service is down on the OMD Insights nodes. Log in to the OMD Insights nodes and enter the command systemctl status splunk to check the status of the Splunk service. The status of the service should show “running”.</p>
service_status for tcp-rabbitmq-exchange	<p>The tcp-rabbitmq-exchange service collects Sensu checks output from the TCP handler and sends the metrics data to the RabbitMQ on the same Monitor node. Failure of this service would result in no metrics data in Influxdb, but Monitor mailer alerts and Uchiwa will still be functional.</p> <p>Connect to the Monitor node and check the status of the process using the command systemctl status tcp-rabbitmq-exchange.</p> <p>Inspect /var/log/messages on the Monitor node for any error messages.</p> <p>Try bringing up the service using sudo systemctl restart tcp-rabbitmq-exchange on the Monitor node. If the alarm persists, then escalate for investigation.</p>
sshd_running_status	<p>Log in to the server through the console and check whether the sshd service is running. If it is not running, start the sshd service and escalate to Level 3.</p>
swap_usage	<p>Try to determine which processes are consuming the swap memory on the system raising the alarm. This issue needs to be escalated ASAP.</p>
tcp_connection	<p>Level 3 should look at the running state of the system to determine if there are any issues.</p> <p>Use the netstat command to check how the connections are distributed.</p>

Check	Remediation
tripwire_running_status	<p>Level 3 should address all issues with these alarms. Level 2 should take no action unless Level 3 determines there is an issue.</p> <p>To clear an alarm, connect to the server and switch to root. Look at the history for the tripwire command. Run the command to set files and then run the command <code>/usr/sbin/tripwire --check</code> to ensure there are no violations.</p>



APPENDIX J

OMD Monitor Alarms and Remediation

This appendix describes items you should check when you triage alerts and describes the checks that are run by the sensu-client service on the OMD Monitor clients, including when alarms or alerts are triggered. It also includes possible remediation steps for any alarms or alerts that are triggered.

When you triage alerts, you should check the following items:

- The active network configuration change requests that could affect traffic
- The active Media Streamer change requests for the servers in question
- Whether the alert indicates that end customer services are impacted, such as the alert is on a Cache node or the OMD Insights server is down
- Run the **dmesg | grep -i error** command and check for any errors
- Check the `/var/log/messages` files for any errors (**grep -i error**)
- Check OMD Insights to check whether the effected node provided content during the error window

This chapter includes the following sections:

- [OMD Monitor Client Checks, on page 309](#)

OMD Monitor Client Checks

The following table describes the checks that are run by the sensu-client service on the OMD Monitor clients. For information on the possible remediation steps you can perform, based on the alarms or alerts that are raised for the checks that the sensu-client service runs, see [Table 14: OMD Monitor Client Alarm Remediation Steps, on page 327](#).



Note

The `/srv/pillar/omd_monitor_checks_config.sls` pillar file on the Salt Master defines the checks that the sensu-client service will perform and their configuration. Based on this file, the `/etc/sensu/conf.d/system_metrics.json` file is created and pushed to each OMD Monitor client. The `/srv/pillar/omd_monitor_checks_config.sls` pillar file contains separate check configurations for both mailer alerts and influxdb metrics. For the mailer alerts, the check names will contain the `"_check"` suffix.



Note To make changes to the /srv/pillar/omd_monitor_checks_config.sls, contact your Cisco Account team.

Table 13: OMD Monitor Client Checks

Check Name	Metrics (check_name in Influxdb)	Entity Family (entity_family in Influxdb)	Metric Value (check_value in Influxdb)	Alarm Condition/ Check Result (check_result in Influxdb) [0 - No Alert/Normal] [1 - Warning Alert] [2 - Critical Alert]	Severity
ats_freelist_metrics	Traffic server freelist statistics	common (will be cache_nodes in future releases)	Reads the Traffic Server freelist and produces trackable output and summarized statistics. Reads from the /opt/trafficserver/var /log/trafficserver/ traffic.out file by default.	No Alert	N/A

Check Name	Metrics (check_name in Influxdb)	Entity Family (entity_family in Influxdb)	Metric Value (check_value in Influxdb)	Alarm Condition/ Check Result (check_result in Influxdb) [0 - No Alert/Normal] [1 - Warning Alert] [2 - Critical Alert]	Severity
balance_verify	balance_verify auto_balance_timer	monitor-node	<p>This check indicates whether the monitor clients are out of balance and if one of the OMD Monitor nodes is over loaded:</p> <ul style="list-style-type: none"> • 0: None of the OMD Monitor nodes are overloaded • 1: One of OMD Monitor nodes is overloaded <p>If balance_verify is 1, the auto_balance_timer value will be a time stamp in the future when the auto load balancing of OMD Monitor clients can be attempted by the OMD Monitor node.</p> <p>Auto load balancing will only happen if all the OMD Monitor nodes agree on the following:</p> <ul style="list-style-type: none"> • auto_balance_timer value • OMD Monitor version • NTP is in sync on all of the OMD Monitor nodes 	Warning	P4
check_bond_interface	bond_interface	cache_nodes	<p>Status of the bond interface:</p> <ul style="list-style-type: none"> • 0: If the bond interface is up and running • 1: If the bond interface is not configured • 2: If the bond interface is down 	Critical: Alert raised if bond interface is configured and down	P4

Check Name	Metrics (check_name in Influxdb)	Entity Family (entity_family in Influxdb)	Metric Value (check_value in Influxdb)	Alarm Condition/ Check Result (check_result in Influxdb) [0 - No Alert/Normal] [1 - Warning Alert] [2 - Critical Alert]	Severity
check_drives	disk_drives_count	cache_nodes	Count of current disk partitions	Critical: Alert raised if the number of disk drive partitions currently available is less than the original number of partitions	P3
check_interfaces	interface	cache_nodes	Status of the slave interfaces of the bond interface: <ul style="list-style-type: none"> • 0: If the interface is up and running • 1: If the interface is not configured • 2: If the interface is down 	Critical: Alert raised if an interface that is part of the bond interface is down	P3
cpu_load_average	<ul style="list-style-type: none"> • Load_avg_1min • Load_avg_5min • Load_avg_15min 	common	CPU load average values, regarding the CPU and I/O over the last 1, 5, and 15 minutes. This value is a decimal number and not percentage. These values are taken from the /proc/loadavg system file.	Warning: <ul style="list-style-type: none"> • Load_avg_1min >= 10 • Load_avg_5min >= 20 • Load_avg_15min >= 30 Critical: <ul style="list-style-type: none"> • Load_avg_1min >= 25 • Load_avg_5min >= 50 • Load_avg_15min >= 75 	Warning: N/A Critical: P4

Check Name	Metrics (check_name in Influxdb)	Entity Family (entity_family in Influxdb)	Metric Value (check_value in Influxdb)	Alarm Condition/ Check Result (check_result in Influxdb) [0 - No Alert/Normal] [1 - Warning Alert] [2 - Critical Alert]	Severity
cpu_usage	cpu_usage	common	CPU usage as a percentage (0 - 100%)	Warning alarm raised if: <ul style="list-style-type: none"> • >= 50% (Default) • >= 80% (Monitor Node) • >= 55% (Edge) Critical alarm raised if: <ul style="list-style-type: none"> • >= 80% (Default) • >= 95% (Monitor Node) • >= 85% (Edge) 	Warning: N/A Critical: P4

Check Name	Metrics (check_name in Influxdb)	Entity Family (entity_family in Influxdb)	Metric Value (check_value in Influxdb)	Alarm Condition/ Check Result (check_result in Influxdb) [0 - No Alert/Normal] [1 - Warning Alert] [2 - Critical Alert]	Severity
disk_metrics	vda_reads vda_readsMerged vda_sectorsRead vda_readTime vda_writes vda_writesMerged vda_sectorsWritten vda_writeTime vda_ioInProgress vda_ioTime vda_ioTimeWeighted	common	I/O statistics of each block device (disk). Each metric is taken from the /proc/diskstats system file, corresponding to the following fields respectively: <ul style="list-style-type: none"> • reads completed successfully • reads merged • sectors read • time spent reading (ms) • writes completed • writes merged • sectors written • time spent writing (ms) • I/Os currently in progress • time spent doing I/Os (ms) • weighted time spent doing I/Os (ms) 	No Alert	N/A
disk_usage	disk_usage	common	Total disk usage for all partitions as a percentage	Warning alarm raised if: <ul style="list-style-type: none"> • >= 85% (Default) • >= 90% (Monitor Node) Critical alarm raised if: <ul style="list-style-type: none"> • >= 95% 	Warning: N/A Critical: P3/P4

Check Name	Metrics (check_name in Influxdb)	Entity Family (entity_family in Influxdb)	Metric Value (check_value in Influxdb)	Alarm Condition/ Check Result (check_result in Influxdb) [0 - No Alert/Normal] [1 - Warning Alert] [2 - Critical Alert]	Severity
dns_status	dns_resolved_status	common	Status of DNS resolution: • 0: DNS is not resolvable • 1: DNS is resolvable	Critical: Alert raised if DNS is not resolvable (check result returns 0)	P4
hardware	Hardware	common	Hardware error message status from the dmesg command output: • 0: No errors found • 1: Errors found	Critical: Hardware error reported in dmesg.	P3
insights_failover	insights_failover_check	splunk	• 0: Failover did not happen and the primary OMD Insights Splunk nodes are active. • 1: Failover happened for the Splunk service on the backup OMD Insights Splunk nodes	Warning: Failover occurred, the standby node is currently the master.	P4

Check Name	Metrics (check_name in Influxdb)	Entity Family (entity_family in Influxdb)	Metric Value (check_value in Influxdb)	Alarm Condition/ Check Result (check_result in Influxdb) [0 - No Alert/Normal] [1 - Warning Alert] [2 - Critical Alert]	Severity
interface_metrics	eth0_rxBytes eth0_rxPackets eth0_rxErrors eth0_rxDrops eth0_rxFifo eth0_rxFrame eth0_rxCompressed eth0_rxMulticast eth0_txBytes eth0_txPackets eth0_txErrors eth0_txDrops eth0_txFifo eth0_txColls eth0_txCarrier eth0_txCompressed	common	<p>Statistics for all configured network interfaces taken from the /proc/net/dev file.</p> <ul style="list-style-type: none"> • Bytes: The total number of bytes of data transmitted or received by the interface. • Packets: The total number of data packets transmitted or received by the interface. • Errs: The total number of transmit or receive errors detected by the device driver. • Drop: The total number of packets dropped by the device driver. • fifo: The number of FIFO buffer errors. • Frame: The number of packet framing errors. • colls: The number of collisions detected on the interface. • Compressed: The number of compressed packets transmitted or received by the device driver. (This appears to be unused in the 2.2.15 kernel.) • Carrier: The number of carrier losses detected by the device driver. • Multicast: The number of multicast frames transmitted or received by the device driver 	No Alert	N/A

Check Name	Metrics (check_name in Influxdb)	Entity Family (entity_family in Influxdb)	Metric Value (check_value in Influxdb)	Alarm Condition/ Check Result (check_result in Influxdb) [0 - No Alert/Normal] [1 - Warning Alert] [2 - Critical Alert]	Severity
N/A	interface_status	caches	<ul style="list-style-type: none"> • 0: The interface is up and running. • 2: The interface is not up and running. 	Critical if any interface on a given client is not up and running	P3
kafka_broker_status	<ul style="list-style-type: none"> • kafka_broker_status_all • kafka_broker_status 	monitor-node Kafka broker IP/hostname	<ul style="list-style-type: none"> • kafka_broker_status_all indicates the number of servers that are having a problem with the Kafka broker server. 0 means all servers are fine or there are no servers configured for Kafka export. • kafka_broker_status indicates the reachability of the Kafka broker server. 0 means it is not reachable. 1 means it is reachable. 	Critical: Any or all Kafka brokers are not able to export metrics.	P3/P4

Check Name	Metrics (check_name in Influxdb)	Entity Family (entity_family in Influxdb)	Metric Value (check_value in Influxdb)	Alarm Condition/ Check Result (check_result in Influxdb) [0 - No Alert/Normal] [1 - Warning Alert] [2 - Critical Alert]	Severity
keepalive	keep_alive	common	<p>The status of the Monitor node connection to the server node being monitored. To determine this, keepalives are sent by the sensu-client service on the server node to a Monitor node every 30 seconds:</p> <ul style="list-style-type: none"> • 0: Not reachable • 1: Reachable 	<p>Warning: Alert raised if no keepalives have been received for 90 to 180 seconds</p> <p>Critical: Alert raised if no keepalives have been received for more than 180 seconds</p>	<p>P3/P2</p> <p>The priority is determined based on the number of servers that are raising the keepalive. If a single server is in this state, it is a P3.</p> <p>If there is a block of servers that are affected, it is a P2 and you should escalate immediately.</p>
login_failed_attempt	Log_Failed_NotAllowed	common	Number of failed attempts made from unauthorized users	<p>Warning: 1 attempt</p> <p>Critical: 2 or more attempts</p>	P4

Check Name	Metrics (check_name in Influxdb)	Entity Family (entity_family in Influxdb)	Metric Value (check_value in Influxdb)	Alarm Condition/ Check Result (check_result in Influxdb) [0 - No Alert/Normal] [1 - Warning Alert] [2 - Critical Alert]	Severity
monitor_client_count	monitor_client_count	monitor-node	The number of clients connected to the OMD Monitor node, including the sensu client running on the OMD Monitor node. This may be -2 for the first time if the sensu server just started on the OMD Monitor node and no clients have connected yet to rabbitmq.	No alert	N/A
monitor_influxdb_memory	monitor_influxdb_memory	monitor-node	<ul style="list-style-type: none"> • 1: Influxdb high on memory and monitor-node is low on available memory. Influxdb memory usage > 70% and system memory > 85%, • 0: Influxdb low on memory or monitor-node has free available memory. 	Warning: Influxdb high on memory and monitor-node is low on available memory. check_result = 1	P4
monitor_version	version_major version_minor version_maint	monitor-node	Monitor version string, separated in to three parts. For example OMD Monitor release 3.9.1 would show: <ul style="list-style-type: none"> • version_major = 3 • version_minor = 9 • version_maint = 1 	No alert	N/A

Check Name	Metrics (check_name in Influxdb)	Entity Family (entity_family in Influxdb)	Metric Value (check_value in Influxdb)	Alarm Condition/ Check Result (check_result in Influxdb) [0 - No Alert/Normal] [1 - Warning Alert] [2 - Critical Alert]	Severity
netstat	netstat	<p>Stored with TCP connection state:</p> <ul style="list-style-type: none"> • ESTABLISHED • SYN_SENT • SYN_RECV • FIN_WAIT1 • FIN_WAIT2 • TIME_WAIT • CLOSE • CLOSE_WAIT • LAST_ACK • LISTEN • CLOSING 	<p>These metrics indicate the number of connections in each state.</p> <p>An alarm is raised if the number of connections exceeds a threshold. A server only has a finite number of TCP connections available, so this alarm indicates when the server is nearing that threshold.</p> <p>Connections are determined by looking for ESTABLISHED and TIME_WAIT entries using the netstat command.</p>	<p>Warning:</p> <p>ESTABLISHED or TIME_WAIT:</p> <ul style="list-style-type: none"> • >= 500 (Default) • >= 40000 (cache nodes) • >= 700 (tcnode) • >= 1500 (monitor node) <p>Critical:</p> <p>ESTABLISHED or TIME_WAIT:</p> <ul style="list-style-type: none"> • >= 1000 (Default) • >= 45000 (cache nodes) • >= 1200 (tcnode) • >= 2000 (monitor node) 	<p>Warning:</p> <p>N/A</p> <p>Critical:</p> <p>P4</p>

Check Name	Metrics (check_name in Influxdb)	Entity Family (entity_family in Influxdb)	Metric Value (check_value in Influxdb)	Alarm Condition/ Check Result (check_result in Influxdb) [0 - No Alert/Normal] [1 - Warning Alert] [2 - Critical Alert]	Severity
ntp_status	<ul style="list-style-type: none"> ntp_offset ntp_sync_status 	common	<ul style="list-style-type: none"> ntp_offset: The offset from the NTP time ntp_sync_status: The status of the sync with the NTP server. 0 indicates the server is not in sync. 1 indicates the server is in sync. 	Warning: ntp_offset >= 100ms Critical: ntp_offset >= 500ms	Warning: N/A Critical: P4
partition_usage	partition_usage	Name of each partition	Disk usage for each partition, as a percentage. This check will return 1(warning) if the configured warning or critical threshold values are more than 100.	Warning: Any partition usage >= 85% Critical: Any partition usage >= 95%	Warning: N/A Critical: P4
password_failed_attempt	Log_Failed_password	common	The number of failed attempts made with incorrect passphrase.	Warning: Alert raised if 1 failed attempt Critical: Alert raised if 2 or more failed attempts	P4
ram_usage	ram_usage	common	RAM usage as a percentage	Warning: Alert raised if RAM Usage is >= 80% Critical: Alert raised if RAM Usage is >= 95%	P4
redis_reset	redis_reset	monitor-node	Indicates whether an OMD Monitor node restarted its Redis instance: <ul style="list-style-type: none"> 0: Not reset 1: Redis is reset 	Warning: If an OMD Monitor node restarted its Redis instance, a warning is raised.	N/A

Check Name	Metrics (check_name in Influxdb)	Entity Family (entity_family in Influxdb)	Metric Value (check_value in Influxdb)	Alarm Condition/ Check Result (check_result in Influxdb) [0 - No Alert/Normal] [1 - Warning Alert] [2 - Critical Alert]	Severity
redis_status	redis_status	monitor-node	The number of connected Redis nodes (OMD Monitor nodes). This number should be equal to the number of OMD Monitor nodes. If it is not, a warning is raised.	Warning: The number of connected Redis nodes is less than the number of OMD Monitor nodes.	N/A
sensu-server is not processing metrics (direct alert from an OMD Monitor node)	—	—	—	Warning: An OMD Monitor node is no longer processing metrics	N/A
service_sshd_status	sshd_running_status	common	The running status of the sshd service: • 0: Not running • 1: Running	Warning: Alert raised if the sshd service is not running.	P4
service_status_grafana-server	service_status	grafana-server	The running status of the Grafana server on the OMD Monitor node: • 0: Not running • 1: Running	Critical: Alert raised when the grafana-server service does not have the status of “running”.	P4
service_status_haproxy	service_status	haproxy	The running status of the haproxy service on the OMD Monitor node: • 0: Not running • 1: Running	Critical: Alert raised when the when the haproxy service does not have the status of “running”.	P4
service_status_influxdb	service_status	influxdb	The running status of the influxdb service on the OMD Monitor node: • 0: Not running • 1: Running	Critical: Alert raised when the when the influxdb service does not have the status of “running”.	P4

Check Name	Metrics (check_name in Influxdb)	Entity Family (entity_family in Influxdb)	Metric Value (check_value in Influxdb)	Alarm Condition/ Check Result (check_result in Influxdb) [0 - No Alert/Normal] [1 - Warning Alert] [2 - Critical Alert]	Severity
service_status_influxdb-relay	service_status	influxdb-relay	The running status of the influxdb-relay service on the OMD Monitor node: • 0: Not running • 1: Running	Critical: Alert raised when the influxdb-relay service does not have the status of “running”.	P4
service_status_rabbitmq-metrics-exporter	service_status	rabbitmq-metrics-exporter	Running status of rabbitmq-metrics-exporter on monitor node • 0: Running • -1: Not Running	Critical when rabbitmq-metrics-exporter service is not in a running state.	P2/P3
service_status_tcp-rabbitmq-exchange	service_status	tcp-rabbitmq-exchange	Running status of tcp-rabbitmq-exchange on monitor node: • 0: Running • -1: Not Running	Critical when tcp-rabbitmq-exchange service is not in a running state.	P2/P3
service_status_rabbitmq-server	service_status	rabbitmq-server	The running status of the rabbitmq-server service on the OMD Monitor node: • 0: Not running • 1: Running	Critical: Alert raised when the rabbitmq-server service does not have the status of “running”.	P4/P3
service_status_redis	service_status	redis	The running status of redis service on the OMD Monitor node: • 0: Not running • 1: Running	Critical: Alert raised when the redis service does not have the status of “running”.	P4/P3

Check Name	Metrics (check_name in Influxdb)	Entity Family (entity_family in Influxdb)	Metric Value (check_value in Influxdb)	Alarm Condition/ Check Result (check_result in Influxdb) [0 - No Alert/Normal] [1 - Warning Alert] [2 - Critical Alert]	Severity
service_status_redis-sentinel	service_status	redis-sentinel	The running status of redis-sentinel service on the OMD Monitor node: • 0: Not running • 1: Running	Critical: Alert raised when the when the redis-sentinel service does not have the status of “running”.	P4/P3
service_status_sensu-server	sensu-server	The running status of sensu-server service on the OMD Monitor node: • 0: Not running • 1: Running	Critical: Alert raised when the when the sensu-server service does not have the status of “running”.	P4/P3	sensu-server
splunk_service	service_status	splunk	The running status of splunk service on the OMD Monitor node: • 0: Not running • 1: Running	Critical: Alert raised when the when the splunk service does not have the status of “running”.	P4/P3
swap_usage	swap_usage	common	The running status of the sshd service: • 0: Not running • 1: Running	Warning: Alert raised if the sshd service is not running.	P2

Check Name	Metrics (check_name in Influxdb)	Entity Family (entity_family in Influxdb)	Metric Value (check_value in Influxdb)	Alarm Condition/ Check Result (check_result in Influxdb) [0 - No Alert/Normal] [1 - Warning Alert] [2 - Critical Alert]	Severity
tcp_connection	tcp_connection	Node type	Number of TCP connections including TCP6	<p>Warning: Alert raised if TCP connections are:</p> <ul style="list-style-type: none"> • >= 500 (others) • >= 50000 (Cache nodes) • >= 5000 (Traffic router) <p>Critical: Alert raised if TCP connections are:</p> <ul style="list-style-type: none"> • >= 1000 (others) • >= 55000 (Cache nodes) • >= 10000 (Traffic router) 	<p>Warning: N/A</p> <p>Critical: P4</p>
tripwire_overrun	tripwire_overrun	common	<p>The status of tripwire overrun:</p> <ul style="list-style-type: none"> • 0: Tripwire is running normally • 1: Tripwire has not completed in a reasonable amount of time 	Critical: tripwire has not completed in a reasonable amount of time	P4

Check Name	Metrics (check_name in Influxdb)	Entity Family (entity_family in Influxdb)	Metric Value (check_value in Influxdb)	Alarm Condition/ Check Result (check_result in Influxdb) [0 - No Alert/Normal] [1 - Warning Alert] [2 - Critical Alert]	Severity
tripwire	tripwire_violations	common	The number of tripwire report violations. Local local files are checked for changes.	Critical: If any tripwire violations are found (check result is 3)	P4
var_partition_usage	var_partition_usage	common	Percentage of /var partition used	Warning: >= 90% Critical: >= 95%	Warning: N/A Critical: P4
vmstat_metrics	procs_waiting procs_uninterruptible memory_swap_used memory_free memory_buffers memory_cache swap_in swap_out io_received io_sent system_interrupts _per_second system_context_ switches_per_second cpu_user cpu_system cpu_idle cpu_waiting cpu_steal	common	Virtual memory statistics. Provides system memory, swap, and processor resource utilization in real time.	No alert	N/A

Table 14: OMD Monitor Client Alarm Remediation Steps

Check	Remediation
ats_freelist_metrics	N/A
balance_verify	<p>Clients are out of balance and one of the OMD Monitor nodes is overloaded. Check the monitor_auto_balance_enable parameter in the /opt/cisco/etc/omd.conf file to see if auto load balancing is enabled.</p> <p>If auto balancing is enabled, this alarm should be resolved automatically. It may take between 5–15 minutes to resolve. If an email confirming that the alarm has been resolved is not received within 30 minutes after the initial alarm email was received, the monitor_auto_balance_client_threshold parameter in the omd.conf file may need to be increased. To avoid overloading any OMD Monitor node, set this value to one of the following:</p> <p>a. $\text{number_of_clients} / \text{number_of_monitor_nodes} \times 1.5$</p> <p>or</p> <p>b. 40</p> <p>Note Each OMD Monitor node can support between 40–50 clients with Kafka metrics export enabled, and up to 70 clients without Kafka metrics export enabled.</p> <p>To manually re-balance a Monitor cluster, perform the following steps from the OMD Salt Master:</p> <ol style="list-style-type: none"> 1. Stop the rabbitmq service: <pre>salt -G 'roles:monitor-node' cmd.run 'systemctl stop rabbitmq-server'</pre> 2. Wait 60 seconds. 3. Start the rabbitmq service: <pre>salt -G 'roles:monitor-node' cmd.run 'systemctl start rabbitmq-server'</pre> 4. Verify that the cluster is operating correctly with: <pre>salt -G 'roles:salt-master' state.apply ail_node_check</pre>
check_bond_interface	<p>If this alarm has been raised, the alarm should have already been cleared so a keepalive alarm should have been raised. Also check the IPMI Interface check in Traffic Ops to see if the server has power.</p>
check_drives	<p>Log in to the cache node and use the <code>ls /dev/sd*</code> command to check the disk drive partitions list.</p> <p>Based on the profile of Mid or Edge cache sever, the list will change. sda, sdb, sdc, and so on are the disk drives and sda1, sda2, and so on are partitions in sda. Check if any disk is missing based on the profile.</p>
check_interfaces	<p>Connect to the server and check the details of the interface that is down.</p>

Check	Remediation
cpu_load_average	<ul style="list-style-type: none"> • Check the running state of the system to determine if there is an issue. If the alert is generated and resolved immediately, it may be able to be ignored. • This may or may not be an issue depending on the load/traffic on the server. If the alert continuously persists, use the top command from the client terminal to determine which process is running high on CPU. • Check the Grafana dashboard for any patterns in the usage.
cpu_usage	<ul style="list-style-type: none"> • Check the running state of the system to determine if there is an issue. If the alert is generated and resolved immediately, it may be able to be ignored. • This may or may not be an issue depending on the load/traffic on the server. If the alert continuously persists, execute the top command on the client terminal to determine which process is running high on CPU. • Check the Grafana dashboard for any patterns in the usage.
disk_metrics	N/A
disk_usage	For critical alerts, log in and check how different partitions are being used. Use the df command to determine which partitions have high disk usage.
dns_status	Make sure servers have access to DNS servers and that those DNS servers are correct.

Check	Remediation
hardware	<p>Indicates a possible hardware error. Details can be found in the <code>/var/log/messages</code> file. The check is looking for “Hardware Error” in the <code>/var/log/messages</code> log file (using the dmesg command). Note that:</p> <ul style="list-style-type: none"> • If “hardware error” is seen in the <code>dmesg</code> output, but has not been corrected in the <code>mcelog</code> (located in <code>/var/log/messages</code>), an alert will be raised immediately. • Alerts are not raised if <code>dmesg</code> lists the same error based on the <code>dmesg</code> time stamp. • Alerts are not raised for corrected errors shown in the <code>mcelog</code> (located in <code>/var/log/messages</code>) • Alerts are raised if the correctable errors are reported more than 6 hours of the day. <p>Perform the following steps:</p> <ol style="list-style-type: none"> 1. If the alert is repeatedly raised, then Stash (mute) the alert while investigating the issue. 2. SSH to the affected system and change to the root user. 3. Collect the relevant portion of the <code>/var/log/messages</code> file for information on the hardware error. 4. Send an email for a case to be opened, including all of the collected information. See the note below. 5. Back up the <code>dmesg</code> log using the following command: dmesg > /var/log/dmesg-`date +%Y%m%d_%I%M%S` 6. Enter the dmesg -c command to clear the <code>dmesg</code> log. 7. Unstash (unmute) the alert. <p>Note Level 1 should create a case with the hardware vendor and ask them to identify what the error means and to recommend the next steps. If the vendor recommends a DIMM replacement, ensure that the issue has actually been identified. Ask the vendor to explain why the DIMM needs to be replaced and to identify which slot is referenced in the error message.</p>
insights_failover	<p>Check the connectivity between the primary and backup OMD Insights nodes. There may be network connectivity issues between the primary and backup OMD Insights nodes that resulted in this alarm.</p> <p>Check for any keepalive alerts from OMD Monitor for the OMD Insights primary nodes.</p>
interface_metrics	N/A
interface_status	Connect to the server and check the details of the interface that is down.

Check	Remediation
kafka_broker_status	<p>Check the accessibility of the Kafka servers from the OMD Monitor nodes:</p> <ul style="list-style-type: none"> • Check the omd.conf pillar file on the Salt Master and verify the Kafka servers and their ports are configured correctly. • SSH to the OMD monitor node and verify connectivity to the Kafka servers. • Check the /var/log/sensu/sensu-server.log and /etc/sensu/kafka_broker_status.log log files for any error messages.
keepalive	<p>Network connectivity between an OMD Monitor node and the server being monitored might have been lost. Additional insight into the network is required.</p> <p>If you can access the server using SSH, you can run the following commands to determine whether the system rebooted or if there was a network issue:</p> <ul style="list-style-type: none"> • uptime: This command shows how long the system has been up. If it has been up for well before the alert, than use the next command to further troubleshoot. • grep “NIC Link” /var/log/messages*: This will display the details for the interface status, if it had a problem. <p>If neither of these commands reveal an issue, perform additional review of the following log files:</p> <ul style="list-style-type: none"> • /var/log/sensu/sensu-client.log on the server • /var/log/sensu/sensu-server.log on Monitor nodes
login_failed_attempt	<p>Level 3: SSH to the server as root and check the /var/log/secure file for a string similar to the following:</p> <pre>"User root from 11.22.33.44 not allowed because not listed in AllowUsers"</pre>
monitor_client_count	<p>This metric can be check to see how the clients are distributed on the OMD Monitor nodes.</p>
monitor_influxdb_memory	<p>When Influxdb is high on memory and the monitor-node is low on available memory, an alert is raised and the Influxdb process is restarted on the monitor-node. This will free memory and the alert will be resolved by itself. After the Influxdb is restarted, it will not be restarted again by this check for 12 hours.</p> <p>If the alarm persists for more than one minute, check the Influxdb service usage using the top command output to see how the memory is being used. Try restarting the Influxdb process by entering the systemctl restart influxdb command and then check the memory usage.</p>
monitor_version	N/A

Check	Remediation
netstat	SSH to the server. Check the output from the netstat command to determine more information about the connections and which port is being used. Look for TCP connections in the ESTABLISHED and TIME_WAIT states, which could provide information about the related service.
ntp_status	Connectivity to the NTP server needs to be checked or other network resources are not responding. Use the ntpq -p command to check the NTP source stratum for correct time.
partition_usage	SSH to the server or view Grafana dashboards for the server to determine which partition is high on usage. After connecting to the server using SSH, look for files that are consuming a large amount of disk space.
password_failed_attempt	SSH to the server as root and check the /var/log/secure file for the string "Failed".
ram_usage	SSH to the server and check the memory usage. Use commands such as top to identify the processes that are running. Check Grafana dashboards for any patterns in the RAM usage.
redis_reset	If this warning is generated continuously, there may be an issue to look at. Check the connectivity between the OMD Monitor nodes; from the salt master run the command salt <salt-master-key> state.apply ail_node_check to verify the cluster status.
redis_status	A warning for this check indicates that an OMD Monitor node could not connect to one or more of the other OMD Monitor node. However, if OMD Monitor is installed in an HA configuration, the OMD Monitor nodes should continue to work without issue. You should investigate the reason for the loss of connection between OMD Monitor nodes. If there is no issue with connectivity, review the /var/log files on the OMD Monitor node.
sensu-server is not processing metrics (direct alert from an OMD Monitor node)	Sometimes this alert will be raised briefly during maintenance windows if a reboot or software upgrade is occurring. At other times, the condition is abnormal; the alert email will continue if the sensu-server service is unable to process metrics. If the alert emails continue, you should investigate the alerting server.
service_sshd_status	Log in to the server through the console and check whether the sshd service is running. If it is not running, start the sshd service.

Check	Remediation
service_status_grafana-server	<p>Connect to the OMD Monitor node and enter the systemctl status grafana-server command to check the status of the grafana-server service. If the service is not in a running state, enter the sudo systemctl restart grafana-server command to try to bring it up.</p> <p>Note If the grafana-server is not in a running state, the Grafana GUI will also fail to load when connected to this OMD Monitor node.</p> <p>Check the /var/log/grafana/grafana.log file for any error messages.</p> <p>If the alarm persists, escalate the problem for further investigation.</p>
service_status_haproxy	<p>Connect to the OMD Monitor node and enter the systemctl status haproxy command to check the status of the haproxy service. If the service is not in a running state, enter the sudo systemctl restart haproxy command to try to bring it up.</p> <p>Check the /var/log/haproxy.log and /var/log/haproxy-status.log files for any error messages.</p> <p>If the alarm persists, escalate the problem for further investigation.</p>
service_status_influxdb	<p>Connect to the OMD Monitor node and enter the systemctl status influxdb command to check the status of the influxdb service.</p> <p>Influxdb metrics may have stopped on one OMD Monitor node, however, the remaining OMD Monitor nodes will continue to store metrics.</p> <p>Enter the sudo systemctl restart influxdb command to try to bring up the influxdb service on OMD Monitor node. If the alarm persists, escalate the problem for investigation.</p>
service_status_influxdb-relay	<p>Connect to the OMD Monitor node and enter the systemctl status influxdb-relay command to check the status of the influxdb-relay service.</p> <p>Influxdb metrics may have stopped on one OMD Monitor node, however, the remaining OMD Monitor nodes will continue to store metrics.</p> <p>Enter the sudo systemctl restart influxdb-relay command to try to bring up the influxdb-relay service on the OMD Monitor node. If the alarm persists, escalate the problem for investigation.</p>

Check	Remediation
service_status_rabbitmq-metrics-exporter	<p>The rabbitmq-metrics-exporter service pops out metrics data from RabbitMQ and exports the same to the Monitor Nodes, InfluxDBs, and the external Kafka servers. Failure of this service would result in no metrics data in Influxdb but the Monitor mailer alerts and Uchiwa will still be functional.</p> <p>Connect to the Monitor node and check the status of the process using <code>systemctl status rabbitmq-metrics-exporter</code>.</p> <p>Inspect <code>/var/log/messages</code> on Monitor node for any error messages.</p> <p>Try bringing up the service using <code>sudo systemctl restart rabbitmq-metrics-exporter</code> on the Monitor node. If the alarm persists, then escalate for investigation.</p>
service_status_rabbitmq-server	<p>Connect to the OMD Monitor node and enter the <code>systemctl status rabbitmq-server</code> command to check the status of the rabbitmq-server service.</p> <p>An OMD Monitor node whose rabbitmq-server service is not running will not be monitoring any clients. Clients previously connected to this node will move to another OMD Monitor node, including the Sensu client of the OMD Monitor node whose rabbitmq-server service is not running.</p> <p>Inspect the <code>/var/log/rabbitmq/</code> logs for any error messages.</p> <p>Enter the <code>sudo systemctl restart rabbitmq-server</code> command to try to bring up the rabbitmq-server service on the OMD Monitor node. If the alarm persists, escalate the problem for investigation.</p>
service_status_redis	<p>Connect to the OMD Monitor node and enter the <code>systemctl status redis</code> command to check the status of the redis service.</p> <p>An OMD Monitor node whose redis service is not running will not be monitoring any clients. Clients previously connected to this node will move to another OMD Monitor node, including the Sensu client of the OMD Monitor node whose redis service is not running.</p> <p>Inspect the <code>/var/log/redis/</code> logs for any error messages.</p> <p>Enter the <code>sudo systemctl restart redis</code> command to try to bring up the redis service on the OMD Monitor node. If the alarm persists, escalate the problem for investigation.</p>
service_status_redis-sentinel	<p>Connect to the OMD Monitor node and enter the <code>systemctl status redis-sentinel</code> command to check the status of the redis-sentinel service.</p> <p>Inspect <code>/var/log/redis/</code> logs for any error messages.</p> <p>Enter the <code>sudo systemctl restart redis-sentinel</code> command to try to bring up the redis-sentinel service on the OMD Monitor node. If the alarm persists, escalate the problem for investigation.</p>

Check	Remediation
service_status_sensu-server	<p>Connect to the OMD Monitor node and enter the systemctl status sensu-server command to check the status of the sensu-server service.</p> <p>Inspect the <code>/var/log/sensu/sensu-server.log</code> file for any error messages.</p> <p>Enter the sudo systemctl restart sensu-server command to try to bring up the sensu-server service on the OMD Monitor node. If the alarm persists, escalate the problem for investigation.</p>
service_status_tcp-rabbitmq-exchange	<p>The tcp-rabbitmq-exchange service collects Sensu checks output from the TCP handler and sends the metrics data to the RabbitMQ on the same Monitor node. Failure of this service would result in no metrics data in Influxdb, but Monitor mailer alerts and Uchiwa will still be functional.</p> <p>Connect to the Monitor node and check the status of the process using the command systemctl status tcp-rabbitmq-exchange.</p> <p>Inspect <code>/var/log/messages</code> on the Monitor node for any error messages.</p> <p>Try bringing up the service using sudo systemctl restart tcp-rabbitmq-exchange on the Monitor node. If the alarm persists, then escalate for investigation.</p>
splunk_service	<p>The Splunk service is down on the OMD Insights nodes. Log in to the OMD Insights nodes and enter the command systemctl status splunk to check the status of the Splunk service. The status of the service should show “running”.</p>
swap_usage	<p>Try to determine which processes are consuming the swap memory on the system raising the alarm. This issue needs to be escalated ASAP.</p>
tcp_connection	<p>Level 3 should look at the running state of the system to determine if there are any issues.</p> <p>Use the netstat command to check how the connections are distributed.</p>
tripwire_overrun	<p>If this alert is raised, there are a number potential causes. Many of the potential causes will require an incident report.</p> <p>Tripwire examines the files on the server hard drive every 15 minutes. Typically, tripwire will complete within about one minute. This alert will be raised if tripwire has been running for 15 minutes.</p> <p>If the server raising this alarm is an active production server, verify whether maintenance is occurring on this server. Some maintenance activities will cause tripwire to run slower than normal, especially if the maintenance involves disk related activity. Otherwise, there may be either a previously undetected disk-related hardware issue on the server or a software failure that has hung the tripwire process. If an immediate cause of this alarm is not apparent, escalate the problem for further analysis to determine if an incident should be raised.</p>

Check	Remediation
tripwire	<p>Connect to the server as root and run the command <code>/usr/sbin/tripwire --check</code> to find any violations. Use the <code>twprint -m r --twrfile /var/lib/tripwire/report/<name>.twr</code> command to print the report to a file. twr files can be found in the <code>/var/lib/tripwire/</code> directory.</p> <p>After performing the integrity check, run the <code>tripwire --update --twrfile /var/lib/tripwire/report/<name>.twr</code> command to update the tripwire database.</p> <p>For more details, refer to the tripwire usage documentation. For example, refer to the documentation available at https://www.centos.org/docs/2/rhl-rg-en-7.2/ch-tripwire.html.</p>
var_partition_usage	<p>The <code>/var</code> partition is used for logging purposes and some services may be logging large amounts of data to this disk.</p> <p>Connect to the server and check which files in the <code>/var</code> partition are filling the disk.</p>
vmstat_metrics	N/A



APPENDIX K

Changing MongoDB Username and Password for OMD Director

This appendix provides the steps to change the username and password used to authenticate to the mongodb, which is used for cdn-config-mgr-db, after OMD Director has been installed.

To change the username and password that is used to authenticate with mongodb, perform the following steps:

1. SSH into the **Repo Server/Salt Master** as `ocdn_adm` and switch to root:
sudo su - root
2. Start the OMD configuration tool:
omd_cfgtool
3. Choose **OMD Director Parameters > Director DATABASE Parameters**. Change the necessary parameters listed in the "Director DATABASE Parameters" table. Also use this table to record the values that you enter for these parameters.

Table 15: Director DATABASE Parameters

Menu Option	Description	omd.conf Parameter	Value
Mongo User	Username for the Mongo user that will have the readWriteAnyDatabase role.	director_database_mongo_user	
Mongo Password	Password for the Mongo user.	director_database_mongo_pw	
Mongo Admin User	Username for the Mongo superuser that will have the root role.	director_database_mongo_admin_user	
Mongo Admin Password	Password for the Mongo superuser.	director_database_mongo_admin_pw	

4. Enter **R** to return to the root menu.

5. Enter **S** to save the settings to the `omd.conf` file and then enter **A** to apply the changes to the underlying pillar files.
6. Enter **E** to exit the OMD Configuration tool menu.
7. Log in to the *primary* OMD Director Master Controller and confirm that the services are operational. The services are operational when they show a status of “Running”:

```
#kubectl get pod
```

NAME	READY	STATUS	RESTARTS	AGE
alarms-2363032894-dxy69	1/1	Running	0	13h
backup-3158271177-4vutp	1/1	Running	0	13h
cdn-config-mgr-1381333575-1lqro	1/1	Running	0	13h
gui-2306176459-wueih	1/1	Running	0	13h
ha-3197191034-ma7hl	1/1	Running	0	13h
kpi-metrics-2549502089-zwhg7	1/1	Running	0	13h
mongodb-3266149902-ydf1x	1/1	Running	0	13h
notification-2332775105-ayhyd	1/1	Running	0	10h
notification-db-2556607229-pnddi	1/1	Running	0	10h
user-management-3177584364-4urfx	2/2	Running	1	13h
user-management-db-3726062850-tx8ng	1/1	Running	0	13h

8. Repeat Step 7 on the *backup* OMD Director Master Controller.



APPENDIX L

Manage Content Invalidation using the OMD Director REST API

This appendix describes how to manage content invalidation jobs using the Cisco Open Media Distribution (OMD) Director REST APIs.

The Cisco OMD Director REST APIs comply with the REpresentational State Transfer (REST) standard. Each URL (called a resource URL) exposes uniform interfaces to the API clients. The API clients use standard HTTP methods of POST, GET, PUT, and DELETE to make calls to the URLs. The HTTP methods are used to describe the create, read, update, and delete (CRUD) actions to be performed.

To manage content invalidation jobs, you will use the following OMD Director REST APIs:

- OMD Director CDN Manager REST API:
 - Is used to access objects that are managed by the OMD Director CDN Manager (cdn-config-mgr) microservice, such as jobs and delivery services.
 - Uses the following URL structure:
`http://<hostname_or_IP_of_OMD_Director_Worker_Node>:8080/<resource-path>`
- OMD Director User Management REST API:
 - Is used to access the objects that are managed by the User Management (microservice) microservice, such as users and authorization tokens.
 - Uses the following URL structure:
`http://<hostname_or_IP_of_OMD_Director_Worker_Node>:8081/<resource-path>`

This chapter includes the following topics:

- [Authentication, on page 339](#)
- [Content Invalidation, on page 341](#)

Authentication

The OMD Director REST APIs use OAuth 2.0 bearer tokens to authorize the API calls for a user. To generate a token for an OMD Director user, use the OMD Director User Management login API call, as described in [Generate Token, on page 340](#). This will authenticate the user and then provide a token that will be used to authorize the API calls. Before you can manage the content invalidation jobs, you must generate a token that

will be used to authenticate the API calls. You will need to include this token in the authorization header of the API request.

**Note**

Make sure you are generating a token for a user that has been assigned a role with privileges to perform the desired task. For example, to add or modify a content invalidation job, the user must have either the Administrator or CDN Admin role. A user with the Administrator, CDN Admin, or CDN Viewer role can view a job or delivery service.

Generate Token

Generates a bearer token to use with the OMD Director API calls that require bearer token authorization.

- Method:

POST

- Request URL:

`http://<hostname_or_IP_of_OMD_Director_Worker_Node>:8081/api/1.0/login`

- Request Header:

"Content-Type: application/json"

- Required Data Parameters:

```
parameters = {
  "username": "<OMD_Director_username>",
  "password": "<OMD_Director_user_password>"
}
```

- Success Response Example:

```
{ "access_token": "MbbGciOiJIUzI1NiIsInR5cCI6IkpXVCJ9.eyJMb2dWQ22JBTEwiLCJjb20uY2lzY28ub21kLmRpcmVjdG9yLnJvbGUiOiM2ZG1pbiIsImV4cCI6MTUxOTk0OTElNiwiXNzIjoiaT01EX1VTRVJNR01UX01TIiwic46iIjoib21kYABtaW4ifQ.ob7q74T9J-H26EjOPSpliImSqObDUyWzySuRqvEC09g", "token_type": "Bearer",
  "name": { "first_name": "user1", "middle_name": "", "last_name": "" }, "email_address": "user1@companyx.example",
  "phone_number": "0", "role": "admin", "status": "reset", "Organization": "", "expiration_date": "0",
  "expiration_warning": "" }
```

- Error Response:

- Invalid:

```
{ "message": "User not authorized, you have X times to retry" }
```

- Invalid user:

```
{ "message": "Invalid user: Username mdbones is not unique/existing!" }
```

- Sample curl Call:

```
curl -d '{"username": "admin1", "password": "abc!23"}' -H "Content-Type: application/json"
http://10.93.22.18:8081/api/1.0/login
```

Content Invalidation

The OMD Director CDN Manager API is used to manage and configure delivery services and content invalidation. The API object for configuring content invalidation is a 'job'.

To configure the content invalidation jobs, you must know the Delivery Service ID (dsId) and Delivery Service name (dsXmlId) for which the content invalidation job will be created. To see a list of all Delivery Services, including their names and IDs, use the OMD Director CDN Manager API `delivery_service` call, described in the List Delivery Services section.


Note

All of the OMD Director CDN Manager API calls use the format of **`http://<hostname_or_IP_of_OMD_Director_Worker_Node>:8080/<resource-path>`**

List Delivery Services

This section discusses the following advanced features of Delivery Services: Returns data for all Delivery Services

- Method:
GET
- Request URL:


Note

The API calls for managing the content invalidation jobs use port 8080, *not* port 8081.

`http://<hostname_or_IP_of_OMD_Director_Worker_Node>:8080/delivery_service/`


Note

To return data for a single Delivery Service use the URL `/delivery_service/<delivery_service_ID>`.

- Request Header:
'Authorization:Bearer <Token>'


Note

`<Token>` is the token you generated for the OMD Director user using the OMD Director User Management API login call, described in the [Generate Token](#) section.

- Success Response Example:

```
{
  "active": true,
```

```

"anonymous_blocking_enabled": false,
"bypass_fqdn": "",
"cache_configuration": "",
"cdn_id": "8",
"content_scope": "national",
"content_type": "vod",
"customer": "",
"device_groups": [],
"display_name": "vod1",
"dns_bypass_cname": "",
"dns_bypass_ip": "",
"dns_bypass_ip6": "",
"dns_bypass_ttl": 30,
"ds_dns_ttl": 3600,
"ds_profile_name": "",
"ds_rfqn": [
  {
    "id": "119",
    "set_number": 0,
    "sub_domain": "vod",
    "type": "SUB_DOMAIN"
  }
],
"ds_rfqn_others": [],
"dscp_mid_tag": -1,
"dscp_tag": 0,
"edge_header_rewrite": "",
"edge_header_rewrite_list": [],
"geo_limit": "None",
"geo_limit_redirect_url": "",
"geo_miss_lat": "41.881944",
"geo_miss_lon": "-87.627778",
"geo_provider": "Maxmind",
"id": "92",
"info_url": "",
"ingest_manifest_proxy": "",
"ingest_manifest_url": "",
"initial_dispersion": 1,
"ipv6_routing": false,
"logs_enabled": false,
"max_bps": "",
"max_dns_answers": "0",
"max_txn_allowed": 0,
"mid_header_rewrite": "",
"mid_header_rewrite_list": [],
"multisite_origin": "no",
"name": "vod1",
"org_server_fqdn": "http://vserver1.cdn.companyx.com",
"origin_shield": null,
"preload_on": 0,
"profile_name": "",
"protocol": [
  "http"
],
"qstring_handler": 0,
"range_request": 0,
"regex_remap": "",
"regional_geo_blocking": false,
"remap_text": "",
"routing_type": "http",
"server_names": [
  "edge-cache4",
  "edge-cache5"
],

```

```

"service": "",
"session_tracking_enabled": false,
"session_tracking_query_key_list": "",
"singed_urls": true,
"signing_algorithm": "url_validator",
"signing_keys_from": "static",
"sixcn_edge": "",
"sixcn_mid": "",
"sixcn_origin": "",
"ssl_key_version": 0,
"status": "success",
"tr_cors_allow_credentials": false,
"tr_cors_allowed_origins": [],
"tr_cors_exposed_headers": [],
"tr_cors_preallowed_headers": [],
"tr_cors_preallowed_methods": [],
"tr_cors_premaxage": "",
"tr_request_headers": "",
"tr_response_headers": "",
"tr_response_headers_list": [],
"url_signing_keys": {},
"use_content_prepositioning": false,
"use_dedicated_volume": ""
}

```



Note In this example, the name of the Delivery Service is vod1 and the Delivery Service ID is 92.

• Sample curl Call:

```

curl -H 'Authorization:Bearer
eyJMbGciOiJIUzI1NiIsInR5cCI6IkpXVCJ9.eyJMbWQiOiJBTEwiLCJjb20uY2lzY28ub21kImRpcmVjdG9yLnJvbGUiOiJy
ZWFKLW9ubHkiLCJleHAiOiJlMjAwMTY0MDYsImZcyI6Ikk9NRF9VU0VSTUdNVF9NUyIsInN1YiI6Imlib25lcYJ9.YT2q6keQ3k
YCD6Wjgl9iu9yNA3-1jtSmP7c2RYhLVuI' http://10.93.22.18:8080/delivery_service/

```

List Content Invalidation Jobs

Returns data for all jobs

• Method:

GET

• Request URL:

http://<hostname_or_IP_of_OMD_Director_Worker_Node>:8080/cdnm/1.1/job/

• Request Header:

'Authorization:Bearer <Token>'



Note <Token> is the token you generated for the OMD Director user using the OMD Director User Management API login call, described in [Generate Token, on page 340](#).

- Success Response Example:

```
{
  "jobs": [
    {
      "dsId": 83,
      "dsXmlId": "vod2",
      "id": "48",
      "jobStatus": "PENDING",
      "operation_id": "dd8283b1-3b1f-4178-9c7d-5ec32cael26c",
      "regex": "http://vod1.companyx.com/live/*.mpd",
      "startTime": "2018-01-17 02:05:33+00",
      "status": "success",
      "ttl": "TTL:54h"
    },
    {
      "dsId": 79,
      "dsXmlId": "vod",
      "id": "47",
      "jobStatus": "PENDING",
      "operation_id": "641f098b-7eec-4340-8daa-f9abb115b3fd",
      "regex": "http://vod2.companyx.com/vod/*.html",
      "startTime": "2018-01-16 23:59:59+00",
      "status": "success",
      "ttl": "TTL:54h"
    }
  ]
}
```

- Sample curl Call:

```
curl -H 'Authorization:Bearer
eyJhbGciOiJIUzI1NiIsInR5cCI6IkpXVCJ9.eyJkdWQiOiJBTEwiLCJjb20uY2l2Y28ub21kLmRpcmVjdG9yLnJvbGUiOiJyZWFrLW9ubHkiLCJleHAiOiE1MjAwMTY0MDYsImZcyI6Ik9NRk9VU0VSTUdNVF9NUyIsInN1YiI6Im1ib25lcyJ9.YT2q6keQ
3kYCD6Wjgl9iu9yNA3-1jtSmp7c2RYhLVuI' http://10.93.22.18:8080/cdnm/1.1/job/
```

List Specific Content Invalidation Job

Returns data for a specific job.

- Method:

GET

- Request URL:

http://<hostname_or_IP_of_OMD_Director_Worker_Node>:8080/cdnm/1.1/job/<id>

- URL Parameters:

Name	Type	Required/Optional	Description
id	Integer	Required	job ID for job to retrieve details for

- Request Header:

'Authorization:Bearer <Token>'



- Success Response Example:

- Sample curl Call:

Create a Content Invalidation Job

- Method

- Request URL:

- Request Header:

- Required Data Parameters:

345

where:

- *<path_to_content_to_invalidate>*: The path to the content that should be invalidated for this Delivery Service. Do not enter the Origin Server Base URL, only enter the path to the content, for example `/video/dec10/*`. The Origin Server Base URL that is configured for the Delivery Service is automatically added to create the complete Regex value.
- *<Delivery_Service_ID>*: The ID of Delivery Service for which the content invalidation job should be created. This value is a positive integer. You can use the CDN Manager API `delivery_service` call, described in [List Delivery Services, on page 341](#), to determine this ID.
- *<#_of_hours_to_run_job>*: The number of hours the content validation rule should be activate. The minimum number is 1. This value is an integer.
- *<job_start_time>*: The date and time at which the content invalidation job should start. The format for this parameter is `YYYY-MM-DD hh:mm:ss`. This must be a day that is between now and two days in the future or you will receive an error message.

• Success Response Example:

```
{ "regex": "http://vserver1.cdn.companyx.com/content/.*", "dsXmlId": "vod1", "jobStatus":
  "PENDING", "startTime": "2018-02-28 11:09:37+00", "ttl": "TTL:1h", "operation_id":
  "3bf4f99f-35fb-4735-8dea-0852967118ce", "dsId": 92, "id": "49" }
```



Note In the output, 'dsXmlId' is the name of the Delivery Service for which the job was created and 'jobStatus' is the status of the job.

• Sample curl Call:

```
curl -X POST -d '{"regex": "/content/.*", "dsId": 92, "ttl": 1, "startTime": "2018-01-11
  18:43:26"}' -H 'Authorization:Bearer
  eyMbbGciOiJIUzI1NiIsInR5cCI6IkpXVCJ9.eyJMbWQzI2JBTEwiLCJjb20uY2IzY28ub21kLmRpcmVjdG9yLnJvbGUiOi
  MbZGlpbiIsImV4cCI6MTUxOTk0OTE1NiwiXNzIjoIT01EX1VTRVJNR01UX01TIiwic46iIjoib21kYABtaW4ifQ.ob7q74T9J-
  H26EjOPSpLiImSqObDUyWzySuRqvEC09g' http://10.93.22.18:8080/cdnm/1.1/job/
```

Delete a Content Invalidation Job

Cancels a content invalidation job.



Note Using the DELETE method with the job API call moves the job to a "CANCELLED" state and does not actually remove it from the list of jobs.

• Method:

DELETE

• Request URL:

`http://<hostname_or_IP_of_OMD_Director_Worker_Node>:8080/cdnm/1.1/job/<jobID>`

• URL Parameters:

Name	Type	Required/Optional	Description
jobID	Integer	Required	ID of job to delete

- Request Header:

'Authorization:Bearer <Token>'



Note <Token> is the token you generated for the OMD Director user using the OMD Director User Management API login call, described in [Generate Token, on page 340](#). To create a job, the user for whom you created the token must have either the Administrator or CDN Admin role.

- Success Response Example:

```
{"operation_id": "1b3013b6-2068-4880-b68e-bc160cc13448"}
```

- Sample curl Call:

```
curl -X DELETE -H 'Authorization:Bearer
eyJMbGciOiJIUzI1NiIsInR5cCI6IkpXVCJ9.eyJModWQiOiJBTEwiLCJjb20uY2l2Y28ub21kImRpcmVjdG9yLnJvbGUiOiM
bZGlpbIIsImV4cCI6MTUxOTkzNTQxMiwiZXNzIjoiaT01EXlVTRVJNR01UX01TIiwic3ViIjoib21kYWRTaW4ifQ.9OWwEYXt
OxK2UZGmEjxcp-Et1ZTqdJCjH89OrsG_vnk' http://10.93.22.18:8080/cdnm/1.1/job/49
```

