

Administrator Guide

CA Application Delivery Analysis

Version 10.1



This Documentation, which includes embedded help systems and electronically distributed materials, (hereinafter referred to as the "Documentation") is for your informational purposes only and is subject to change or withdrawal by CA at any time.

This Documentation may not be copied, transferred, reproduced, disclosed, modified or duplicated, in whole or in part, without the prior written consent of CA. This Documentation is confidential and proprietary information of CA and may not be disclosed by you or used for any purpose other than as may be permitted in (i) a separate agreement between you and CA governing your use of the CA software to which the Documentation relates; or (ii) a separate confidentiality agreement between you and CA.

Notwithstanding the foregoing, if you are a licensed user of the software product(s) addressed in the Documentation, you may print or otherwise make available a reasonable number of copies of the Documentation for internal use by you and your employees in connection with that software, provided that all CA copyright notices and legends are affixed to each reproduced copy.

The right to print or otherwise make available copies of the Documentation is limited to the period during which the applicable license for such software remains in full force and effect. Should the license terminate for any reason, it is your responsibility to certify in writing to CA that all copies and partial copies of the Documentation have been returned to CA or destroyed.

TO THE EXTENT PERMITTED BY APPLICABLE LAW, CA PROVIDES THIS DOCUMENTATION "AS IS" WITHOUT WARRANTY OF ANY KIND, INCLUDING WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT. IN NO EVENT WILL CA BE LIABLE TO YOU OR ANY THIRD PARTY FOR ANY LOSS OR DAMAGE, DIRECT OR INDIRECT, FROM THE USE OF THIS DOCUMENTATION, INCLUDING WITHOUT LIMITATION, LOST PROFITS, LOST INVESTMENT, BUSINESS INTERRUPTION, GOODWILL, OR LOST DATA, EVEN IF CA IS EXPRESSLY ADVISED IN ADVANCE OF THE POSSIBILITY OF SUCH LOSS OR DAMAGE.

The use of any software product referenced in the Documentation is governed by the applicable license agreement and such license agreement is not modified in any way by the terms of this notice.

The manufacturer of this Documentation is CA.

Provided with "Restricted Rights." Use, duplication or disclosure by the United States Government is subject to the restrictions set forth in FAR Sections 12.212, 52.227-14, and 52.227-19(c)(1) - (2) and DFARS Section 252.227-7014(b)(3), as applicable, or their successors.

Copyright © 2014 CA. All rights reserved. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

CA Technologies Product References

This document references the following CA Technologies products:

- CA Application Delivery Analysis
- CA GigaStor
- CA GigaStor Connector
- CA Multi-Port Monitor
- CA Performance Center
- CA NetQoS Performance Center
- CA Standard Monitor
- CA Virtual Systems Monitor

Contact CA Technologies

Contact CA Support

For your convenience, CA Technologies provides one site where you can access the information that you need for your Home Office, Small Business, and Enterprise CA Technologies products. At <http://ca.com/support>, you can access the following resources:

- Online and telephone contact information for technical assistance and customer services
- Information about user communities and forums
- Product and documentation downloads
- CA Support policies and guidelines
- Other helpful resources appropriate for your product

Providing Feedback About Product Documentation

If you have comments or questions about CA Technologies product documentation, you can send a message to techpubs@ca.com.

To provide feedback about CA Technologies product documentation, complete our short customer survey which is available on the CA Support website at <http://ca.com/docs>.

Contents

Chapter 1: Welcome to Application Delivery Analysis 17

Understanding Application Delivery Analysis.....	17
Architectural Overview of Application Delivery Analysis	18
Monitoring with Application Delivery Analysis	19
The Administrator Role in Application Delivery Analysis	20
How to Access the Management Console.....	20
Recommended Browser Settings	21
How to Navigate the Administration Page.....	22
Monitored TCP Sessions.....	23
View the List of Client Networks	23
View the List of Servers.....	24
View the List of Applications.....	26
How to Set Up and Maintain CA Application Delivery Analysis	27

Chapter 2: Managing Client Networks 29

How Client Networks Work.....	29
How the Network List Works	30
How Network-Based Reporting Works	31
How Network Regions Work	32
Naming Conventions	34
Find a Client Network.....	34
Manage Client Networks.....	37
Default Client Networks	38
Import Client Networks from a CSV File.....	39
Add a Client Network	44
Edit a Client Network	46
Delete a Client Network	47
SNMP Poll a Router for Client Networks.....	48
Export Client Networks to a CSV File.....	49
Group Client Networks by Network Type	49
How Network Types Work	50
Why Network Types Are Useful	51
Add a Network Type.....	52
Edit a Network Type.....	53
Delete a Network Type	53
Assign a Network Type to a Client Network.....	54

Manage Client Networks with Web Service Methods	54
Parameter Descriptions	54
Web Service Methods	56
How to Test the Web Services API	60
How Error Reporting Works	61
Sample Perl Script	62

Chapter 3: Managing Servers 67

How Servers Work.....	67
How the Server Subnet List Works.....	68
How the Server List Works.....	69
How /32 Client Networks Work	70
TCP Session Identification	71
Host Name Resolution	71
Manage Server Subnets	72
Add a Server Subnet.....	72
Edit a Server Subnet.....	74
Delete a Server Subnet.....	75
Manage Servers.....	76
Naming Conventions	77
Find a Server	77
Add a Server	78
Edit a Server	79
Delete a Server.....	80
Import Server Definitions from a CSV File.....	81
Export Server Definitions to a CSV File.....	85
Pinning a Monitor Feed to a Server	86
Schedule Server Maintenance.....	87
How Maintenance Schedules Work	88
Add a Maintenance Schedule	89
Rename a Server Maintenance Schedule	90
Delete a Maintenance Schedule	90
Add a Maintenance Period to a Maintenance Schedule.....	91
Edit a Maintenance Period.....	92
Delete a Maintenance Period	92
Assign a Maintenance Schedule to a Server	93

Chapter 4: Managing Tenants 95

Introduction to Tenancy.....	95
Prerequisites	96

How Domains Separate Traffic.....	97
Data Source Synchronization	97
How Domain-Based Reporting Works.....	98
Add Client Networks to a Domain	98
Add Servers to a Domain.....	99
Assign Domains to Monitor Feeds	100

Chapter 5: Managing Applications **103**

How Applications Work.....	103
How Priority Applications Work	104
Find an Application	105
Naming Conventions	106
Application Port Exclusions	106
How Port Exclusions Work	107
How the Port Exclusion List Works	108
Add a Port Exclusion.....	109
Edit a Port Exclusion.....	111
Delete a Port Exclusion	112
Manage System-Defined Applications	113
Edit a System-Defined Application.....	114
Delete a System-Defined Application	115
Manage User-Defined Applications	116
Create a Standard Application	117
Create a Web Application	119
Create an FTP Application	123
Create a Control Port Application	125
Assign Servers to an Application	126
Edit a User-Defined Application	128
Delete a User-Defined Application.....	129
Manage a Multi-Tiered Application	130
How Multi-Tiered Applications Work.....	131
How To Monitor a Multi-Tiered Application	132
Application Keep-Alive Messages.....	136

Chapter 6: Managing Performance Thresholds **137**

How Performance Thresholds Work	138
How Application Performance is Rated	140
How Performance Metrics Work	141
Options to Customize Performance Thresholds	144
How Incidents Open and Close	147

NetQoS Performance Center (CA NPC)	148
CA Performance Center (CA PC).....	149
Edit Performance Thresholds	150
Edit Thresholds from the Administration Page	151
Edit Thresholds from the Operations Page	152
Add Performance Thresholds.....	154
Enable Default Performance Thresholds for a Group of Networks.....	155
Edit Performance Thresholds for WAN-Optimized Network Segments.....	156
Edit Thresholds for Non-Optimized Traffic on an Optimized Application.....	157
Edit Thresholds for the Optimized Client Segment	158
Edit Thresholds for the Optimized WAN Segment.....	160
Edit Thresholds for the Optimized Server Segment.....	162

Chapter 7: Managing Incident Responses 165

How Incident Responses Work	165
How an Incident Response is Launched	166
Email Notifications	168
SNMP Trap Notifications	169
Application Connection Time Investigations.....	170
Packet Capture Investigations.....	171
Performance via SNMP Investigations	172
Ping Response Time Investigations	173
Trace Route Investigations	174
Add an Incident Response	175
Edit an Incident Response	176
Delete an Incident Response.....	177
Add an Action to a Network or Server Incident Response	177
Edit a Responsive Action	178
Delete a Responsive Action	178
Assign an Incident Response	179
Assign an Incident Response to an Application.....	180
Assign an Incident Response to a Server.....	180
Assign an Incident Response to a Network Type	181
Troubleshoot Incident Responses	181
Manage Incidents Using Web Service Methods.....	181
Object Identifier Specifications	182
Web Service Specifications	184
Interpret SNMP Traps	187

Chapter 8: Managing Application Performance OLAs **189**

How Performance OLAs Work.....	189
How Performance OLA Reporting Works.....	190
How Performance OLA Thresholds Work	191
How Operational Level Metrics Work	192
Performance OLA Tips and Tricks.....	193
Establish Operational Levels from Historical Data	194
Create an Application Performance OLA for a Group of Networks	196
Edit an Application Performance OLA	198
Delete an Application Performance OLA.....	199

Chapter 9: Managing Application Availability **201**

How Availability Monitoring Works	201
Why System-Defined Applications are Excluded	202
How Server Incidents for Application Availability Work	203
How Application Availability Reporting Works	203
Enable Availability Monitoring.....	204
How Application Availability OLAs Work.....	207
How Availability OLA Reporting Works	207
Enable an Application Availability OLA	208

Chapter 10: Managing User Account Permissions **209**

How User Account Permissions Work.....	210
Integrated Security.....	211
Product Privileges	212
Roles.....	212
Users and Groups.....	214
Frequently Asked Questions	217

Chapter 11: System Administration **219**

Windows Administrator Credentials	219
Manage the Database	219
Required Services.....	220
The Status of the Database	221
Edit Database Storage Preferences.....	221
Purge Data from the Database.....	222
Back Up and Restore the Database.....	223
Manage Console Settings	223
Change the IP Address	224

Manage SNMP Profiles	225
How SNMP Profile Discovery Works	226
Add a SNMP Profile	227
Edit a SNMP Profile	228
Delete a SNMP Profile	228
Manage Network Devices	229
Add a Network Device	229
View Network Device Investigations	230
Edit a Network Device	230
Delete a Network Device	231
Group Network Devices for Investigation	231
Manage Scheduled Email	233
Edit Schedules for Email Reports	233
Delete a Scheduled Report	234
Perform System Maintenance	234
How to Maintain Hard Disk Drives	234
How to Update System Security and Install Windows Updates	235
Ensure Data Integrity and Use Anti-Virus Software	236
Issues with Third-Party Software	236
Issues with Domain Group Policies	237
Product Upgrade Support	237
Request a Hardware Replacement	237

Chapter 12: Monitoring Device Administration 239

How Monitoring Devices Work	239
How Monitor Feeds Work	240
How Monitor Feed Assignment Works	241
How Monitoring Device Synchronization Works	242
Create a Pair of Monitor Feeds	243
View Sessions Information	244
View Active Sessions on a Monitor Feed	245
View an Hourly Summary of Session Counts	246
How the Management Console Manages Database Growth	246
Database Capacity	247
Database Growth Control	248
Monitoring Device Ratios	249
Monitoring Device Recommendations	251
Perform Basic Operations	252
Manage Monitoring Device Incidents	253
View Monitoring Device Incidents	254
Edit Monitoring Device Incident Thresholds	255

Enable and Disable Availability Monitoring	256
Add an Incident Response to a Monitoring Device	257
Edit Monitoring Device Incident Response Name.....	258
Delete a Monitoring Device Incident Response	259
Add an Action to a Monitoring Device Incident Response.....	259
Edit a Responsive Action	260
Delete a Responsive Action.....	260
Assign a Monitoring Device Incident Response	261
Troubleshoot Monitoring	261
View the Monitoring Device Status.....	262
Troubleshoot Communication Issues.....	263
Troubleshoot Missing Data	264

Chapter 13: Monitoring with a CA Standard Monitor 265

How a CA Standard Monitor Works as a Monitoring Device	265
How the CA Standard Monitor Works	266
Required Services	267
How Monitor Feeds Work	268
How Monitor Feed Assignment Works	268
How Packet Capture Investigations Work.....	269
Monitoring Device Considerations.....	269
Support for XFF Translation.....	270
How XFF Translation Works	270
Enable XFF Translation	271
Add a CA Standard Monitor	272
Prerequisites	272
Add a CA Standard Monitor	273
NAT Firewall Communication.....	274
Secure Packet Capture Investigation Files	275
Edit a CA Standard Monitor.....	276
Edit the Packets Monitor Feed	277
Manage Monitoring Device Performance	278
Monitoring Device Operations	279
Filter Out Keep-Alive Messages.....	281
Delete a CA Standard Monitor	283
Disable the Packets Monitor Feed	284
Troubleshoot a CA Standard Monitor	285
Verify Active Sessions.....	286
View Monitor Feed Statistics	287
View SPAN Receiver Statistics.....	289
Troubleshoot the CA ADA Monitor Service.....	291

Check for Duplicate Client Networks	297
Troubleshoot Missing Data	298
Troubleshoot Dropped Packets.....	299

Chapter 14: Monitoring with a CA Virtual Systems Monitor 301

How a CA Virtual Systems Monitor Works as a Monitoring Device	302
Plan for Deployment	303
Port Usage and Firewalls.....	304
System Requirements	305
Add a CA Virtual Systems Monitor	306
Configure the Virtual Switch	306
Create the Virtual Machine	312
Configure the Network Connections.....	314
Run the CA Application Delivery Analysis Setup Program	317
Synchronize Time with a Time Server	318
Finish the Setup.....	318
Post-Installation Steps.....	319

Chapter 15: Monitoring with a CA GigaStor 321

How a CA GigaStor Works as a Monitoring Device	321
How the CA GigaStor Connector Works.....	322
How Monitor Feed Assignment Works	323
How Packet Capture Investigations Work.....	323
Sizing Recommendations	323
Monitoring Device Considerations.....	324
Add a CA GigaStor Monitoring Device.....	325
Prerequisites	325
Install and Configure Software on the GigaStor Appliance	326
Add a CA GigaStor Monitoring Device	327
Assign a CA GigaStor to a Monitoring Device.....	328
Install CA Observer on the User's Computer	330
Give the User Permission to a Passive Probe Instance	331
Block a CA GigaStor Input Port.....	332
Edit a CA GigaStor Monitoring Device.....	333
Edit the GigaStor Monitor Feed	334
Unassign a CA GigaStor	335
GigaStor Incidents	335
Perform Basic Operations	336
Delete a CA GigaStor Monitoring Device	336
Delete a CA GigaStor	337

Troubleshoot a CA GigaStor Monitoring Device	337
View Active Sessions on the GigaStor Monitor Feed	338
View GigaStor Counter Statistics.....	339

Chapter 16: Monitoring with Cisco WAAS 341

How Cisco WAAS Works as a Monitoring Device	342
How Cisco WAAS Works.....	343
How Monitor Feed Assignment Works	344
How Network Segments Work.....	344
How Performance Thresholds Work for WAN-Optimized Network Segments.....	345
How Monitoring Works when Optimization Stops	345
Sizing Recommendations	347
Monitoring Device Considerations.....	348
Add a Cisco WAE Monitoring Device.....	349
Prerequisites	349
Configure the Cisco WAE to Export Response Time Data	350
Assign a Cisco WAE to a Monitoring Device.....	351
Edit a Cisco WAE Monitoring Device	352
Edit the WAN Opt Monitor Feed	353
Unassign a Cisco WAE	354
WAAS Incidents	354
Delete a Cisco WAE Monitoring Device.....	355
Disable Flow Monitoring on a Cisco WAE	355
Delete a Cisco WAE Monitoring Device	356
Reset Optimized Applications	356
Troubleshoot a Cisco WAE Monitoring Device.....	357
View Active Sessions	357
View WAN Opt Counter Statistics	358
Verify the Cisco WAE Configuration.....	360
Monitor a Server with a Group of Cisco WAE Devices	361
How Source Sets Work.....	361
Assign a Source Set to a Cisco WAE Device.....	362
Assign a Source Set to a Server	363
Rename a Source Set	363
Delete a Source Set	364
Share Optimization Data between Management Consoles	365
Share WAN-Optimized Performance Data.....	367
Update the Shared Configuration	368
Delete a Monitoring Device	369
Troubleshooting Tips.....	370

Chapter 17: Monitoring with Cisco NAM 371

How a Cisco NAM Works as a Monitoring Device	371
How a Cisco NAM Works.....	372
How Monitor Feed Assignment Works	372
Monitoring Device Considerations.....	373
Add a Cisco NAM Monitoring Device	374
Prerequisites	374
Configure a Cisco NAM to Export Response Time Data	375
Verify the Cisco NAM is Connected to the management console	376
Enable the NAM Monitor Feed	377
Edit a Cisco NAM Monitoring Device	378
Edit the NAM Monitor Feed	379
NAM Incidents.....	380
Delete a Cisco NAM Monitoring Device	381
Troubleshoot a Cisco NAM Monitoring Device	383

Chapter 18: Monitoring with Riverbed Steelhead 387

Concepts.....	387
Introduction	387
Architecture	388
Network Segments.....	389
Monitor Feed Assignment.....	389
Incident Thresholds for Network Segments.....	390
Packet Capture Investigations.....	390
Monitoring when Optimization Stops.....	390
Add a Monitoring Device.....	392
Monitoring Device Considerations.....	393
Provision a Monitoring Device	394
Mirror Network Traffic.....	394
Add the Monitoring Device	399
Manage a Monitoring Device	400
Edit a Monitoring Device.....	400
Edit a Monitor Feed	401
Manage Monitor Performance	402
Delete a Monitoring Device	402
Troubleshoot a Monitoring Device	403
View Steelhead Receiver Statistics.....	404
View SPAN Receiver Statistics	406
View Active Sessions	407

Glossary

409

Index

427

Chapter 1: Welcome to Application Delivery Analysis

This section contains the following topics:

[Understanding Application Delivery Analysis](#) (see page 17)

[How to Access the Management Console](#) (see page 20)

[Monitored TCP Sessions](#) (see page 23)


[How to Set Up and Maintain CA Application Delivery Analysis](#) (see page 27)

Understanding Application Delivery Analysis

CA Application Delivery Analysis is the application performance management module of the CA Performance Center (CA PC) and the CA NetQoS Performance Center (CA NPC). CA Application Delivery Analysis tracks and measures end-user response time--without desktop or server agents. CA Application Delivery Analysis passively monitors IPv4-based TCP packets traversing the network between clients and servers, providing metrics such as network, server, and application latency for all mission critical applications.

When a TCP transaction flows across your infrastructure, it essentially flows through three main components of your infrastructure – the network, the server, and the application. When performance degrades on any of these components, it can adversely affect the transaction times to the end users.

CA Application Delivery Analysis measures TCP transaction time from the server switch, between each TCP client and the application port on a server. As shown in the example below, response times between the server switch and the client are network response times (indicated with a green arrow), and response times between the server switch and the server are server response times (indicated with a red arrow). The combined network and server response times (indicated with a yellow arrow) reflect the end-user's experience with the application.



CA Application Delivery Analysis reports on the unique combination of time frame, application port, server, network, and performance metric. For example, you can report on the average Network Connection Time for all applications and servers that communicated with the Development I client network in the last 24 hours.

```
Timeframe: Last 24 Hours
Application: All
Server: All
Network: Development I (192.168.8.0/24) [Clear]
Metric: Network Connection Time
```

CA Application Delivery Analysis can be registered as a data source with the CA PC or the CA NPC when it is used with performance management. When CA Application Delivery Analysis is not used with performance management, it should be registered with the CA NPC. The CA PC or CA NPC security features—including users, roles, product privileges, and groups—let you control which users can view specific data in the management console.

Architectural Overview of Application Delivery Analysis

CA Application Delivery Analysis consists of a CA Application Delivery Analysis Manager and one or more monitoring devices. The CA Application Delivery Analysis Manager is an appliance that runs the database engine and reporting console. A monitoring device is an appliance that monitors TCP transactions. The CA Application Delivery Analysis Manager consolidates response time data from one or more monitoring devices.

CA Application Delivery Analysis offers multiple monitoring options:

- The CA Multi-Port Monitor is the most scalable and feature-rich option for SPAN or TAP data. It has single Gigabit and 10-Gigabit options.
- The CA Standard Monitor processes up to a single Gigabit of SPAN or TAP data.
- The CA Virtual Systems Monitor monitors to-from traffic between virtual servers on the same VMware ESX Host.
- Cisco Network Analysis Module (NAM). A Cisco NAM is a monitoring device for CA Application Delivery Analysis and lets you perform deep-dive troubleshooting.

A CA Standard Monitor or CA Multi-Port Monitor also processes packet digests from:

- WAN-optimization devices, including Cisco WAAS and Riverbed Steelhead
- CA GigaStor Connector. A CA GigaStor is a long-term packet capture appliance that sends performance data to CA Application Delivery Analysis, but also enables retrospective network analysis to go back in time for packet analysis or to replay session data.



In environments where all the server traffic can be mirrored from a single switch port, a Standalone CA Application Delivery Analysis Manager can be used where the CA Application Delivery Analysis Manager receives mirrored TCP packets on its Monitor NIC.

Monitoring with Application Delivery Analysis

CA Application Delivery Analysis monitors end-to-end performance by automatically:

- Collecting and calculating IPv4-based TCP transaction times for applications, based on default server subnets and client networks, and the server subnets and client networks you specify.
- Monitoring user-defined applications, based on the server assignments you specify.
- Generating baselines for application performance to understand what is normal.
- Setting thresholds for upper boundaries of acceptable performance
- Creating incidents when the thresholds are crossed.
- Displaying 5-minute summaries of all performance metrics for IPv4 transactions.

By default, CA Application Delivery Analysis passively monitors end-to-end response time. You can make the product more active, for example, by enabling a network trace route investigation in response to a network incident or setting an operational level agreement (OLA) for application availability.

More information:

[How to Set Up and Maintain CA Application Delivery Analysis](#) (see page 27)

The Administrator Role in Application Delivery Analysis

The CA Application Delivery Analysis administrator is responsible for:

- Setting up monitoring devices.
- Working with network administrators to place monitoring devices at the appropriate locations on the network, and mirror the appropriate TCP traffic to each monitoring device.
- Specifying the server subnet, application, and client network definitions that you want to monitor.
- Managing the applications that the management console automatically creates from all the observed server traffic, for example, to assign an incident response, or to create an application and assign particular servers.
- Specifying levels of service for performance and availability.
- Managing security to enable, for example, users to log in to the management console and manage a particular group of applications, servers, and client networks. Note that when registered as a data source with the CA PC or CA NPC, these tasks are performed in the Admin tab of the CA PC or CA NPC management console.

How to Access the Management Console

After the CA Application Delivery Analysis Manager is installed and available on your enterprise network, open a web browser and type in the IP address of the server that hosts the CA Application Delivery Analysis Manager. If DNS is configured to resolve the host name, enter the host name of the server.

Note For information about installing and configuring CA Application Delivery Analysis, see the *Installation Guide*.

Access the management console web interface by providing the appropriate URL:

If the CA PC is	Specify this URL
Installed on the management console server	<code>http://hostname/sa</code>

If the CA PC is	Specify this URL
Not installed on the management console server	<code>http://hostname</code>

The management console was designed for display in Microsoft Internet Explorer version 8 or 9 and Mozilla Firefox, and requires Adobe Flash Player. Internet Explorer versions 6 and 7 are not supported. If you see a security prompt about blocked websites from Internet Explorer Enhanced Security when you first try to access the Web interface, we recommend you adjust your browser settings.

If you do not have a CA Application Delivery Analysis administrator account, contact your CA PC or CA NPC administrator. If you are logging on for the first time, you can use the predefined administrator account, admin, with the default password, admin.

On the Login page, type your user name and password. Keep in mind that login credentials are case-sensitive. As soon as you have been authenticated successfully, you are automatically redirected to the management console. An incorrect login returns you to the Login page.

For better security, we recommend changing the default passwords for the predefined user accounts. The administrator can do this by logging in and editing user accounts.

More information:

[Managing User Account Permissions](#) (see page 209)

[Recommended Browser Settings](#) (see page 21)

Recommended Browser Settings

We recommend updating your web browser configuration to disable pop-up blocking and, on Internet Explorer only, add the management console to the list of trusted internet sites.

When accessing the management console through Internet Explorer 8, the page formatting appears incorrect at the top of the page.

To avoid this issue, press F12 in Internet Explorer and then set the Browser Mode to IE8 and the Document Mode to Quirks.

Disable Pop-Up Blocking

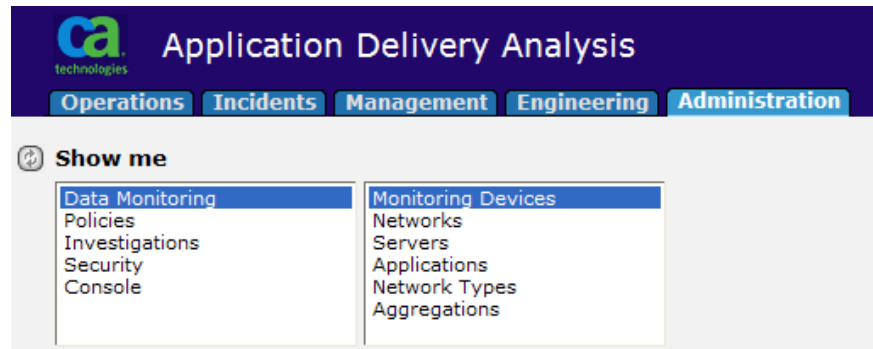
When launching an investigation manually, disable pop-up blockers on the web browser, otherwise the investigation will be blocked from running. If you run an investigation and do not see a pop-up showing the status of the investigation as it runs, then your pop-up blocker is most likely still enabled, and your investigation did not run.

Update Trusted Internet Sites

We recommend adding the host name of the management console server to the list of trusted internet sites in the Internet Explorer browser instance to improve user interface performance. By default, Internet Explorer uses high security settings that restrict navigation to trusted sites or repeatedly display a warning message when you navigate to sites that are not on the list of trusted sites. This recommendation is only applicable to Internet Explorer.

How to Navigate the Administration Page

The Administration page is accessible to CA Application Delivery Analysis users with administrator privileges.



Click the Show Me menu items on the Administration page to manage CA Application Delivery Analysis:

- **Data Monitoring**--Specifies monitoring devices that collect application response time data from the applications, servers, and networks in your environment.
- **Policies**--Specifies thresholds for application, server, and network performance, responses to performance degradation, and operational level agreements (OLAs) for performance and availability.
- **Investigations**--Specifies the SNMP profiles that are used to poll a server or network device.
- **Security**--Specifies user permissions. When registered as a data source with the CA PC or CA NPC, user permissions are managed from the CA PC or CA NPC.
- **Console**--Manages the CA Application Delivery Analysis Manager database, reporting, and monitoring device incidents.

Monitored TCP Sessions

To enable the management console to track and measure end-user response time, a monitoring device must observe an IPv4-based TCP transaction between a client and an application port on a server. Using the server subnets and client networks you specify, the management console automatically calculates performance metrics for matching TCP transactions.

Typically, a CA Technical Representative helps you specify the client networks, servers, and applications that you want to monitor.

If you are managing an ISP with overlapping server or client IP addresses, you will need to use domains to separate the TCP traffic.

More information:


[Managing Tenants](#) (see page 95)

View the List of Client Networks

Use the Network List to manage the list of client networks that the management console monitors. Each monitoring device computes response time metrics between the application servers and the client networks specified in the Network List.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Networks in the Show Me menu.



The Network List appears. In the following example, two IPv4 client networks with a /24 subnet mask are defined.

A screenshot of a 'Network List' table. The table has a header with columns: Network, Subnet, Regions, Network Type, and a set of action icons. Below the header, there are two rows of data. The first row is 'Development I' with subnet '192.168.8.0/24', region '1', and network type 'LAN'. The second row is 'Development II' with subnet '192.168.9.0/24', region '1', and network type 'LAN'. The table also includes a grey header with instructions and a footer with pagination and size controls.

Network	Subnet	Regions	Network Type	
Development I	192.168.8.0/24	1	LAN	<input type="checkbox"/>
Development II	192.168.9.0/24	1	LAN	<input type="checkbox"/>

CA Application Delivery Analysis includes the following default entries for catch all default client networks:

- All Others 0.0.0.0/0
- All Other 10 Networks 10.0.0.0/8
- All Other 172.16 Networks 172.16.0.0/12
- All Other 192.168 Networks 192.168.0.0/16

More information:

[Managing Client Networks](#) (see page 29)


View the List of Servers

Use the Server List to view and manage the list of servers that the management console monitors.

If a server is not displayed in the list, browse the Server Subnets List to verify that a matching server subnet exists. The management console automatically adds a server to the Server List after it observes matching application port traffic on the server.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Servers in the Show Me menu.



The Server List appears. In the following example, the list of IPv4 addresses in the Server List corresponds to the IPv4 server subnets in the Server Subnet List.

[Go to Server List](#)

Server Subnet List							
View and manage server subnets to specify the range of server IP addresses monitored by SuperAgent. When matching traffic on a server is observed, the server is added to the Server List.							
<input type="button" value="Add Server Subnet"/>							
Description	Subnet	Exclusions	Apps	Servers			
CA Client Network	130.200.39.0/24	0	1	1	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
My_subnet	10.0.38.0/24	0	0	0	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
NQ Client Network	192.168.0.0/24	0	13	7	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

[Go to Server Subnet List](#)

Server List								
View and manage the servers that are monitored by SuperAgent. If you do not see a particular server, go to the Server Subnet List and verify that a matching server subnet exists.								
<input type="button" value="Add Server"/>		<input type="text"/>		<input type="button" value="Search"/>		<input type="button" value="Clear"/>		
Server	Address	Collection Device	Apps	Bytes	User Modified	Last Seen		
130.200.39.244	130.200.39.244	Stndalone32-W01 - Primary SPAN Feed	<input type="checkbox"/>	0	0	Yes	Inactive	<input type="checkbox"/>
192.168.0.2	192.168.0.2	Stndalone32-W01 - Primary SPAN Feed	<input type="checkbox"/>	0	0	No	Inactive	<input type="checkbox"/>
192.168.0.20	192.168.0.20	Stndalone32-W01 - Primary SPAN Feed	<input type="checkbox"/>	1	1.8M	No	Current	<input type="checkbox"/>

Note: Application Delivery Analysis includes a default server subnet called *Monitor All Servers* on subnet 0.0.0.0 mask 0. The default server subnet enables automatic data collection to occur when a new monitor is added.

More information:

[Managing Servers](#) (see page 67)


View the List of Applications

Use the Application List to manage the applications that the management console monitors. CA Application Delivery Analysis automatically monitors all application traffic between the specified server subnets and client networks.

If an application is not displayed in the list, make sure the management console is configured to monitor the application.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Applications in the Show Me menu.



The Application List appears. In the following example, application traffic on well-known ports is automatically named.

Application List

View and manage the applications that are monitored by SuperAgent. The Application List includes system-defined applications that are automatically created from observed application port traffic and user-defined applications.

Search [] Clear Search Reset List

Show: Subnets Servers Show: All Applications Max per Page: 10

Create New Application Edit Delete SLAs Thresholds

Application	TCP Port	Type	Servers	Subnets	Bytes	Last Seen	Configured By
<input type="checkbox"/> America Online	5190	Standard	19	1	1.6G	2010-09-23 09:02 CDT	System
<input type="checkbox"/> BSD Remote Login	513	Standard	19	1	170.4M	2010-09-30 10:20 CDT	System
<input type="checkbox"/> Domain Name Service Protocol	53	Standard	19	1	267.6M	2010-09-30 10:20 CDT	System
<input type="checkbox"/> Hypertext Transfer Protocol	80	Standard	14	1	138.9G	2010-09-23 09:05 CDT	System
<input type="checkbox"/> Microsoft SQL Server	1433	Standard	18	1	1.7G	2010-09-30 10:20 CDT	System
<input type="checkbox"/> NETBIOS Session Service	139	Standard	19	1	57.9G	2010-09-30 10:20 CDT	System
<input type="checkbox"/> Port 5060	5060	Standard	19	1	866.6M	2010-09-30 10:20 CDT	System
<input type="checkbox"/> Secure Shell	22	Standard	19	1	8.9G	2010-09-30 10:20 CDT	System
<input type="checkbox"/> Simple Mail Transfer Protocol	25	Standard	19	1	1.4T	2010-09-30 10:20 CDT	System

Create New Application Edit Delete SLAs Thresholds Max per Page: 10

More information:

[Managing Applications](#) (see page 103)

How to Set Up and Maintain CA Application Delivery Analysis

Now that you are familiar with how CA Application Delivery Analysis works, you are ready to perform common administrator tasks.

Tasks	More Information
<ul style="list-style-type: none"> ■ Add or remove a monitoring device ■ Manage monitoring device performance ■ View monitoring device activity 	Monitoring Device Administration (see page 239)
<ul style="list-style-type: none"> ■ Specify the client networks you want to monitor ■ Group networks using network types 	Managing Client Networks (see page 29)
<ul style="list-style-type: none"> ■ Specify the servers you want to monitor ■ Specify periods when server maintenance is scheduled to occur 	Managing Servers (see page 67)
<ul style="list-style-type: none"> ■ Specify the applications you want to monitor 	Managing Applications (see page 103)
<ul style="list-style-type: none"> ■ Add default performance thresholds to groups of networks ■ Customize performance thresholds for a user-defined application across a group of networks ■ Adjust performance thresholds to make CA Application Delivery Analysis more or less sensitive to changes in application, server, and network response time 	Managing Performance Thresholds (see page 137)
<ul style="list-style-type: none"> ■ Enable CA Application Delivery Analysis to automatically notify you and investigate a server or network incident 	Managing Incident Responses (see page 165)
<ul style="list-style-type: none"> ■ Specify operational levels for application performance 	Managing Application Performance OLAs (see page 189)
<ul style="list-style-type: none"> ■ Monitor application availability and specify operational levels for availability 	Managing Application Availability (see page 201)
<ul style="list-style-type: none"> ■ Add or remove users ■ Manage user permissions 	Managing User Account Permissions (see page 209)

Tasks	More Information
<ul style="list-style-type: none"><li data-bbox="488 331 967 394">■ Schedule and perform weekly database backups<li data-bbox="488 415 834 443">■ Specify system preferences<li data-bbox="488 464 786 491">■ Manage SNMP profiles	Managing the Console (see page 219)
<ul style="list-style-type: none"><li data-bbox="488 512 857 539">■ Perform system maintenance	Performing System Maintenance (see page 234)

Chapter 2: Managing Client Networks

This section contains the following topics:

[How Client Networks Work](#) (see page 29)

[Manage Client Networks](#) (see page 37)

[Group Client Networks by Network Type](#) (see page 49)

[Manage Client Networks with Web Service Methods](#) (see page 54)

How Client Networks Work

A *client network* specifies a range of client IPv4 addresses that you want to monitor and corresponds to a physical location or group of users in your environment.

We recommend creating a client network based on where the client traffic originates, such as a remote site, the data center, or for an externally facing application (the Internet). For example, if you have a client network in Austin that is defined as 192.168.0.0/22, use the management console to create an Austin client network with the same network address and subnet mask.

To enable a management console user to quickly analyze and respond to performance issues on a client network, create client networks that correspond to the actual client networks in your environment.

When a server is added to the Server List, the management console automatically creates a corresponding /32 client network to monitor traffic between servers in a multi-tier application.

Properly specifying the client networks of interest optimizes available system resources.

Application Delivery Analysis includes the following default entries for catch all default client networks:

- All Others 0.0.0.0/0
- All Other 10 Networks 10.0.0.0/8
- All Other 172.16 Networks 172.16.0.0/12
- All Other 192.168 Networks 192.168.0.0/16

Note: When client networks overlap, the management console reports the application traffic in the more specific client network.

More information:

[How /32 Client Networks Work](#) (see page 70)

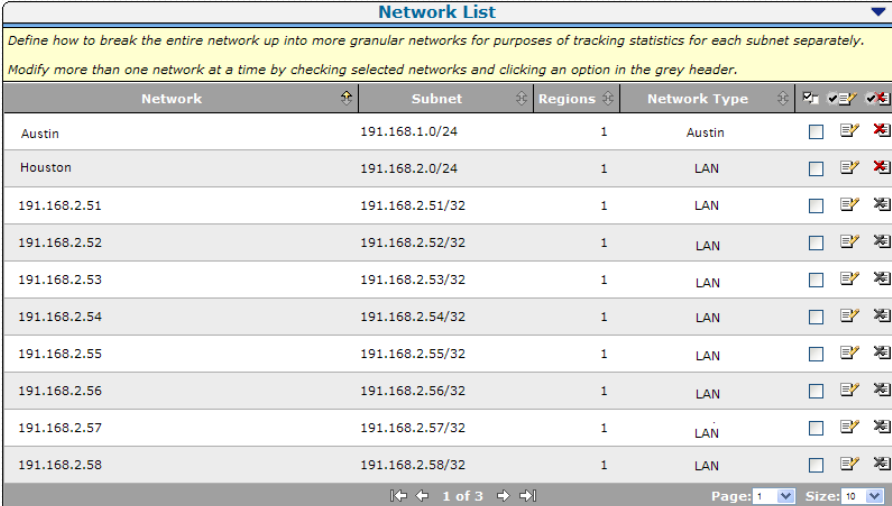
[How the Management Console Manages Database Growth](#) (see page 246)

How the Network List Works

Use the Network List to view and manage the networks that are monitored by the management console. The Network List includes the client networks you specify, plus /32 client networks.

A /32 client network is a client network with a single IPv4 address. CA Application Delivery Analysis uses /32 client networks to monitor multi-tier applications. When a server is added to the Server List, the management console automatically creates a corresponding /32 client network.

Note: To delete a /32 client network, delete the corresponding server from the Server list. You cannot delete a /32 client network that was created by the management console.



Network	Subnet	Regions	Network Type	
Austin	191.168.1.0/24	1	Austin	<input type="checkbox"/>
Houston	191.168.2.0/24	1	LAN	<input type="checkbox"/>
191.168.2.51	191.168.2.51/32	1	LAN	<input type="checkbox"/>
191.168.2.52	191.168.2.52/32	1	LAN	<input type="checkbox"/>
191.168.2.53	191.168.2.53/32	1	LAN	<input type="checkbox"/>
191.168.2.54	191.168.2.54/32	1	LAN	<input type="checkbox"/>
191.168.2.55	191.168.2.55/32	1	LAN	<input type="checkbox"/>
191.168.2.56	191.168.2.56/32	1	LAN	<input type="checkbox"/>
191.168.2.57	191.168.2.57/32	1	LAN	<input type="checkbox"/>
191.168.2.58	191.168.2.58/32	1	LAN	<input type="checkbox"/>

More information:

[Manage a Multi-Tiered Application](#) (see page 130)

How Network-Based Reporting Works

To monitor application performance across the network, the management console reports the average application response times across a client network rather than reporting on the response times for individual TCP sessions. Reporting at the network-level enables the management console to isolate a performance problem to a particular network.

When monitoring a client network that is defined with at least a 24 bit subnet mask (24 bit to 32 bit), the management console reports the actual client IP addresses that communicate with an application during a specified time interval. For example, reporting on the 191.168.1.0/24 network displays the actual client IP addresses that accessed the application.

Because the management console averages the application response time across the network, this information does not help you isolate which client computers were affected by a performance issue. Rather, this is simply a list of clients that communicated with the application during the specified time period. This level of detail is not required to isolate the network as the source of a performance problem.

Address - Mask Listings for Users			
Address	Mask	Domain	Host
191.168.1.0	255.255.255.255	Default Domain	(Not available)
191.168.1.1	255.255.255.255	Default Domain	(Not available)
191.168.1.2	255.255.255.255	Default Domain	(Not available)
191.168.1.3	255.255.255.255	Default Domain	(Not available)
191.168.1.4	255.255.255.255	Default Domain	(Not available)
191.168.1.5	255.255.255.255	Default Domain	(Not available)
191.168.1.6	255.255.255.255	Default Domain	(Not available)

If you create a client network with a subnet mask that is less than 24 bits, the management console reports the Class C networks that communicate with the application. For example, if you create a client network with a 22 bit subnet mask, such as 192.168.0.0/22, the management console reports metrics from the Class C client networks, 191.168.0.0/24 through 191.168.3.0/24, instead of the actual client IPs that accessed the application.

Address - Mask Listings for Users			
Address	Mask	Domain	Host
191.168.0.0	255.255.255.0	Default Domain	(Subnet)
191.168.1.0	255.255.255.0	Default Domain	(Subnet)
191.168.2.0	255.255.255.0	Default Domain	(Subnet)
191.168.3.0	255.255.255.0	Default Domain	(Subnet)

Note: If you are interested in session-level reporting, use a CA Multi-Port Monitor as the monitoring device. Unlike a CA Standard Monitor, which analyzes TCP sessions at the network level with a 5-minute granularity, a CA Multi-Port Monitor can analyze a TCP session between a server and a particular client at a 1-minute granularity. In addition, a CA Multi-Port Monitor lets you analyze traffic volume for all observed network traffic, including TCP and non-TCP traffic.

How Network Regions Work

To enable a management console user to quickly analyze and respond to performance issues at a remote site, create client networks that correspond to the actual client networks in your environment.

If you want to enable the management console to report the actual client IPv4-based IP addresses that communicate with an application, the client networks must have at least a 24 bit subnet mask.


If necessary, you can convert a large client network into smaller network regions so that the management console can report the observed TCP traffic by region. A *network region* is a smaller subnet of a client network. For example, if you have /22 client network, convert it into four /24 network regions.

If the network is defined as a	Create this many regions
/24	1
/23	2
/22	4
/21	8
/20	16
/19	32
/18	64
/17	128
/16	256

When you convert a client network into regions, the management console reports the observed TCP traffic by region, rather than by the defined client network. Continuing the previous example, after you convert the /22 client network into /24 network regions, the /22 client network is only visible from the Administration page and is not available for reporting purposes. In reports, you would need to look for the individual network regions instead.

To enable a management console user to report on application performance across a defined client network rather than on the observed traffic for a particular /24 network region, use the CA PC or the CA NPC to create a group of all the network regions.

To easily identify the network regions that correspond to a client network, use a familiar naming convention. In the example below, the Austin group includes the /24 network regions that are converted from the 192.168.0.0/22 client network.



When reporting on a group of network regions, the performance detail reports on the Engineering page, such as the Response Time Composition: Average report, aggregate data from all the /24 network regions. However, all other reports, such as the Operations page reports and the Engineering page performance maps, list each /24 network region separately. For information about reporting on a group of networks, see the *User Guide*.

The following table shows some sample network configurations with network regions.

Network Configuration	/24 Network Regions
Name: ABC Subnet: 10.10.1.0 Mask: 24 Regions: 1 Network Type: 128K	10.10.1.0
Name: DEF Subnet: 10.10.0.0 Mask: 22 Regions: 4 Network Type: 128K	10.10.0.0 10.10.1.0 10.10.2.0 10.10.3.0

Network Configuration	/24 Network Regions
Name: XYZ	10.10.0.0
Subnet: 10.10.0.0	10.10.1.0
Mask: 16	.
Regions: 256	.
Network Type: 128K	.
	10.10.255.0

More information:

[Naming Conventions](#) (see page 34)

Naming Conventions

Here are some suggested naming conventions for client networks. Keep in mind, if you have defined a client network using regions, you cannot report on the client network. So, from a reporting perspective, the client network region names are important:

For a	Use	For example
Client network	<i>Location-Description</i>	Singapore-Backbone
Client network region	<i>Location-Region-Bandwidth</i>	Austin-Subnet10-128K

Find a Client Network

Search the Network List to find client networks that are currently monitored by the management console. Alternatively, you can use the QoS Users report on the Engineering page to find client IPs that are communicating with an application, for example, on the catch all default client network.


More information:

[Default Client Networks](#) (see page 38)

Search the Network List

Search the Network list to find a client network that is already defined.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Networks in the Show Me menu.
3. (Optional) Choose the domain where you want to search.
4. Click the blue gear menu () in the Network List and select Find Network.
5. Click Search in the Networks list to search by Network Name or Subnet.
6. In the Search For field, specify a search string. You can specify a pattern-matching character at the beginning or end of the search pattern, including:

*

Matches all characters.

%

Matches one character.

7. Click Search to find matching networks.

Note: Click a hyperlink to edit the properties of a client network.

More information:


[Managing Tenants](#) (see page 95)

[Edit a Client Network](#) (see page 46)

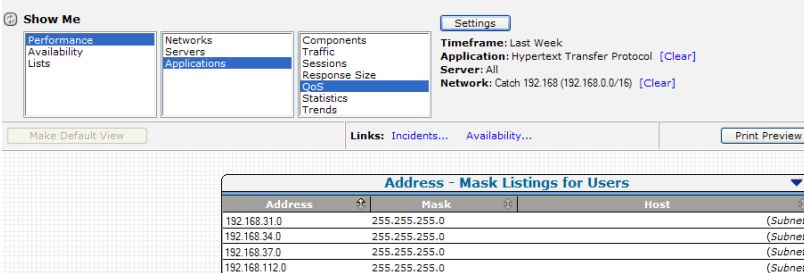
Find Unknown Client Networks with the QoS Users Report

It is common for address ranges to expand over time. Creating a "catch all" network ensures that once that expansion occurs CA Application Delivery Analysis will have calculations for it. If you see client IP traffic show up in the broad superset network, use the QoS Users report to see the /24 client networks that are showing traffic and then explicitly define those subset networks.

Follow these steps:

1. Click the Engineering page.
2. Click the blue gear menu () in the QoS Users report and click List Users to filter the observed traffic on the client network into smaller subnets that are 8 bits more specific. For example, when reporting on a /16 network, such as a 192.168.0.0/16 "catch all" network, the Users list displays the corresponding /24 networks where the management console has observed client IP traffic. In the example below, the /24 networks are expected to be in the range of 30-39, but the management console sees unexpected client IPs in the 192.168.112.0/24 client network.

To see the actual client IPs in the 192.168.112.0/24 client network, create a 192.168.112.0/24 client network and use the QoS Users report to list the IP addresses of the /32 users.



The screenshot shows the QoS Users report interface. On the left, there are navigation menus for Performance, Availability, Lists, Networks, Servers, Applications, Components, Traffic, Sessions, Response Size, QoS, Statistics, and Trends. The main area displays a table titled "Address - Mask Listings for Users" with columns for Address, Mask, and Host. The table lists four entries: 192.168.31.0, 192.168.34.0, 192.168.37.0, and 192.168.112.0, all with a mask of 255.255.255.0 and labeled as (Subnet). The interface also includes a "Settings" button, a "Timeframe" dropdown set to "Last Week", an "Application" dropdown set to "Hypertext Transfer Protocol", a "Server" dropdown set to "All", and a "Network" dropdown set to "Catch 192.168 (192.168.0.0/16)".

Address	Mask	Host
192.168.31.0	255.255.255.0	(Subnet)
192.168.34.0	255.255.255.0	(Subnet)
192.168.37.0	255.255.255.0	(Subnet)
192.168.112.0	255.255.255.0	(Subnet)

CA Application Delivery Analysis includes the following default entries for catch all default client networks:

- All Others 0.0.0.0/0
- All Other 10 Networks 10.0.0.0/8
- All Other 172.16 Networks 172.16.0.0/12
- All Other 192.168 Networks 192.168.0.0/16

More information:

[Default Client Networks](#) (see page 38)

Manage Client Networks

Manage client networks by specifying the networks where you want to monitor application port traffic:

- Import network information from a CSV file to speed the configuration.
- Add a client network.
- Import network information by polling a router.

When adding client networks to the management console, plan to group your client networks by network type for reporting purposes and for launching incident responses on a network.

More information:

[Group Client Networks by Network Type](#) (see page 49)

Default Client Networks

It is common for address ranges to expand over time. A catch all network ensures that once that expansion occurs CA Application Delivery Analysis has calculations for it. If you see the client IP traffic show up in a broad default client network, use the QoS Users report to see the /24 client networks that are showing traffic and explicitly define those subset networks.

CA Application Delivery Analysis includes the following default entries for catch all default client networks:

- All Others 0.0.0.0/0
- All Other 10 Networks 10.0.0.0/8
- All Other 172.16 Networks 172.16.0.0/12
- All Other 192.168 Networks 192.168.0.0/16

Note: When client networks overlap, the management console reports the application traffic in the more specific client network.

Applications that face the Internet in the DMZ can be monitored with the All Others client network for subnet 0.0.0.0, mask 0 to monitor application activity for client networks that you have not explicitly defined. Do not use this client network to monitor other types of applications. If you do, the Operations and Engineering pages will show many unusual client IP addresses that are over threshold because they have no normal baseline.

Note: The All Others client network does not let you see TCP sessions, but does let the management console report application activity from other client networks.

More information:

[How Network-Based Reporting Works](#) (see page 31)

[Find a Client Network](#) (see page 34)

Import Client Networks from a CSV File

Import IPv4 client networks from a CSV file to easily specify the client networks you want.

To get started, you export a list of IPv4 client networks, for example, from a DHCP management tool, and then edit the list to describe the /24 client networks.

To avoid editing your client networks after you import them, we recommend defining all the network properties for each network in the CSV file. Note that after you import a network, you must edit the network from the management console and you cannot re-import an existing network to update its properties.

When creating your CSV file, do not omit network types. Network types play an important role in monitoring application performance.

Note: If you have defined domains in the CA PC or the CA NPC, you must import the list of client networks into the same domain. The supported CSV file syntax does not let you designate a domain for each network definition.

More information:

[Edit a Client Network](#) (see page 46)

[Group Client Networks by Network Type](#) (see page 49)

Create a CSV File

Follow the steps below to create a CSV file that specifies the IPv4 client networks that you want to monitor.

When creating a CSV file:

- Separate each field with a comma and ensure no spaces follow the comma.
- Enclose embedded commas or double quotes inside double quotes.
- Enclose strings that contain one or more spaces in double quotes; for example: "Houston Office".

Follow these steps:

1. Create a file with a .csv file extension.
2. Add each network definition on a separate line.
3. For each network definition, use the following format:

network_name,subnet,mask,regions,network_type

Note that in the example below, the Houston Office entry does not specify a network type, so it defaults to Unassigned in the management console:

"Atlanta Lab",192.168.100.0,24,1,LAN

"Houston Office",192.168.200.0,24,1

network_name

Defines the name the actual client network, up to 50 characters. We recommend using a naming convention, for example, *Location-Description*.

If you are planning to use network regions to convert the network subnet into /24 subnets, we recommend using a naming convention, for example, *Location-Region-Bandwidth*.

subnet

Defines the IPv4 address of the network in four-part, dotted-decimal notation, for example, 192.168.100.0.

mask

Specifies a subnet mask for the actual client network.

Note that the management console automatically creates a /32 network for each server to monitor servers that act as both a server to client requests and a client to another server.

regions

(Optional) Specifies the number of /24 subnets that can be converted from the broader network subnet. For example, if the network subnet is /22, four regions would convert the /22 into four /24 subnets.

network_type

(Optional) Specifies the network type you want to assign to the network and if applicable, its network regions. If you do not specify a network type, it is set to Unassigned.

4. Save the file and then import it into the management console.

More information:

[Manage a Multi-Tiered Application](#) (see page 130)


[How Network Regions Work](#) (see page 32)

[Group Client Networks by Network Type](#) (see page 49)


Import a CSV File

After you create a CSV file, import IPv4 client network definitions into the management console.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Networks in the Show Me menu.
3. (Optional) Choose the domain into which you want to import the client networks.
The Network List opens.
4. Click the blue gear menu () and select Import from File.
Import Network Definitions opens.
5. Click Browse to select the CSV file.
6. Click Next.
Network Properties opens.
7. Resolve any overlap between the definitions you imported from the CSV file and the existing network definitions in the management console. Any overlap is highlighted:

Red

Indicates that the subnet definition matches a subnet definition that is already defined. Click the red x () to delete the network or modify its subnet definition.

You cannot add the same subnet definition more than once. For example, if an 11.2.0.0/16 network with 256 client regions is already defined, you cannot import an 11.2.3.0/24 client network.

Yellow

Indicates the subnet definition overlaps with an existing, broader subnet definition. If you no longer need the broader subnet definition, remove it after you complete the import.

The management console imports network definitions that overlap each other, however, the management console always reports client activity in the more specific network. For example, if you import a /26 network that overlaps with a /24 network, the management console reports matching client traffic in the /26 network.

Green

Indicates the subnet definition overlaps with an existing, narrower subnet definition. If you no longer need the need the narrower subnet definition, remove it after you complete the import.

The management console imports network definitions that overlap each other, however, the management console always reports client activity in the more specific network. For example, if you import a /26 network that overlaps with a /24 network, the management console reports matching client traffic in the /26 network.

8. Click OK.
9. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

More information:

[Managing Tenants](#) (see page 95)

Add a Client Network

Add a client network to enable the management console to monitor the performance of all application servers across a client network. The client network is not assigned to a particular application server.

To quickly and easily import all of your client networks, import network definitions from a CSV file.

To manually add a client network,


Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Networks in the Show Me menu.
3. (Optional) Choose the domain where you want to create the client network.
4. Under the Show Me menu, click to add an IPv4 client network.
5. Complete the fields in the Network Properties and click OK.

For information about network properties, click Help.

6. Resolve any overlap between the definitions you imported from the CSV file and the existing network definitions in the management console. Any overlap is highlighted:

Red

Indicates the subnet definition matches a subnet definition that is already defined. Click the red X () to delete the network or modify its subnet definition.

You cannot replace an existing subnet definition. For example, if an 11.2.0.0/16 network with 256 client regions is already defined, you cannot import an 11.2.3.0/24 client network.

Yellow

Indicates the subnet definition overlaps with an existing, broader subnet definition. If you no longer need the broader subnet definition, remove it after you complete the import.

The management console imports network definitions that overlap each other; however, the management console always reports client activity in the more specific network. For example, if you import a /26 network that overlaps with a /24 network, the management console reports matching client traffic in the /26 network.

Green

Indicates the subnet definition overlaps with an existing, narrower subnet definition. If you no longer need the narrower subnet definition, remove it after you complete the import.

The management console imports network definitions that overlap each other, however, the management console always reports client activity in the more specific network. For example, if you import a /26 network that overlaps with a /24 network, the management console reports matching client traffic in the /26 network.

7. Click OK.
8. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

More information:

[Managing Tenants](#) (see page 95)

[Import Client Networks from a CSV File](#) (see page 39)

Edit a Client Network

Use the Network List to edit a client network or one of its regions. You cannot change the IP address or subnet mask after the management console observes traffic on the network.

Note that you cannot change the domain to which a network is assigned. If necessary, delete a client network and then recreate the client network with the correct domain.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Networks in the Show Me menu.
The Network List appears.
3. (Optional) Choose the domain that contains the client network you want to edit.
4. Click the edit icon (✎) to edit a network.
5. Complete the fields in the Network Properties and click OK.
For information about network properties, click Help.
6. (Optional) To edit a particular network region:
 - a. Click + to expand the list of regions for a client network.
 - b. In the Network List, rename a network region or assign a network type.
 - c. Click OK.
7. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

More information:

[Managing Tenants](#) (see page 95)


Delete a Client Network

Delete a client network to stop monitoring TCP sessions on the client network. Existing network data continues to be available for reporting purposes.

When deleting a client network, keep in mind that you cannot delete:

- A network region from a client network. When you delete a client network, all of its network regions are deleted.
- A /32 or /128 client network that was automatically created by the management console. Delete the corresponding server from the Server list.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Networks in the Show Me menu.
The Network List opens.
3. (Optional) Choose the domain that contains the client network you want to delete.
4. In the Network List, click the red X () to delete a network.
5. In the Delete Confirmation prompt, click Continue with Delete.
6. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

More information:

[Managing Tenants](#) (see page 95)

[How /32 Client Networks Work](#) (see page 70)


SNMP Poll a Router for Client Networks

You can SNMP poll a router for its IPv4 network definitions and import them into the management console. Note that while the management console polls a router, you cannot perform other tasks in the management console.

When querying a router for network information, try to import the directly connected networks that have a route cost of zero. For example, poll a single router at a remote site and retrieve all networks with LAN type costs (for example, less than 5 or 10 or 18) instead of polling 6 to 7 routers at the one site. Using Max Cost limits the networks retrieved to only LOCAL networks so you can associate them to the particular remote site, for example, by following the standard naming convention and grouping the networks using a network type.

To enable the management console to SNMP poll a router, add a network device for the router with an assigned SNMP profile. If you do not assign a SNMP profile to the device, the management console attempts to discover a valid SNMP profile from the list of available SNMP profiles.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Networks in the Show Me menu.
3. (Optional) Choose the domain into which you want to import the client networks.
4. In the Network List, click the blue gear menu () and select Import from SNMP. Import Network Definitions opens.
5. Complete the fields and click Poll.

For information about network definition properties, click Help.

The Poll Confirmation message notifies you that polling might take minutes. Note that while the management console polls the router, you cannot perform other tasks in the management console.

6. Click Continue at the Poll Confirmation.
7. Click Close Status Window after the management console finishes polling to view the results.
8. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

More information:


[Managing Tenants](#) (see page 95)

[Manage Network Devices](#) (see page 229)

Export Client Networks to a CSV File

Export your existing client networks from the management console to a CSV file, for example, to generate a formatted list of networks that you can edit and then import into the management console. Note that the management console does not let you import an existing client network.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Networks in the Show Me menu.
The Network List opens.
3. (Optional) Choose the domain from which you want to export the client networks.
4. Click the blue gear menu () and select Export Networks.
5. Choose an option in the File Download dialog box:
 - To open the file in Microsoft Excel, click Open.
 - To save the file, click Save. In the Save As dialog box, choose a directory and file name for the file and click Save.

More information:

[Managing Tenants](#) (see page 95)

Group Client Networks by Network Type

Network types let you group and manage client networks with similar latency characteristics. For example, you can group the client networks at the Cary site, and customize how the management console manages performance at that location.

By default, CA Application Delivery Analysis includes pre-configured VPN and Wireless network types. Performance thresholds for VPN and Wireless network types are set to one-half the sensitivity of other network types due to increased latency.

How Network Types Work

A *network type* is essentially a group of networks. The group should reflect the latency experienced by the users on the client network.

In general, the largest contributing factors to latency should be the geographic location and the bandwidth to the subnet. The management console includes default network types to indicate bandwidth, such as T1, however, the management console cannot assume where your networks are located relative to the monitoring devices that observes the application traffic. Therefore, it is important to customize network types to your organization's layout because a big performance difference exists between a T1 that is 200 miles long and a T1 that is 1,500 miles long.

The number of network types required depends on the size of your organization. Use network types to group subnets that must talk across the same network paths and therefore experience the same latency. The key principal is that network types should be a group of subnets that share the same physical network resources and therefore experience the same latency due to distance, serialization, and queuing delay. Having an accurate network diagram helps greatly in building network types.

Note: By default, CA Application Delivery Analysis includes pre-configured VPN and Wireless network types for users coming from networks with highly variable latency. Performance thresholds for VPN and Wireless network types are pre-configured to one-half the sensitivity of other network types. As the administrator, you want the management console to create incidents that your team can resolve. In this case, it unlikely you can resolve latency issues with remote users accessing the network from the internet, for example, from a hotel in London or Singapore.

More information:

[Edit Performance Thresholds](#) (see page 150)

Why Network Types Are Useful

Network types provide several benefits:

- Automatically group the client networks for reporting. When registered as a data source with the CA PC or the CA NPC, each network type you define creates a Network Region Type system group.

If you do not assign network types to your client networks, you can create groups in the CA PC or the CA NPC to report on a group of client networks in the management console and the CA PC or the CA NPC.

Consider implementing rule-based site groups in the CA PC or the CA NPC to manage remote sites.

- Tune performance thresholds by network type and enable some client networks to have tighter or looser thresholds on the same application.

The management console calculates sensitivity thresholds for all the TCP sessions between each application port, server, and client network. However, with network types, you can adjust the sensitivity of an application on a particular group of networks.

If you are planning to use static thresholds, network types let you set static thresholds on a particular group of networks.

- Customize network incident responses by network type. Some examples are listed below:
 - Notify the appropriate person by email or SNMP trap.
 - Launch a Trace Route investigation in response to a network incident on a particular group of networks.
- Create performance Operational Level Agreements (OLAs) by network type to allow remote and local networks to have separate performance OLA thresholds. Using network types to group networks with similar latencies ensures that you can properly configure OLAs on Network Round Trip Time and Total Transaction Time metrics by letting you configure your OLAs per network type.

More information:

[Managing Performance Thresholds](#) (see page 137)

[Managing Application Performance OLAs](#) (see page 189)

Add a Network Type

Add a network type to identify a group of client networks that share common latency characteristics, such as a group of client networks at a remote site which share the same physical path to the domain controller.

Name a network type using the same name as the remote site that contains the client networks you want to group. When naming a network type, keep in mind that the management console cannot determine the network path that TCP packets take between the client and server. For example, if you want to create a network type for the Subnet 10 and Subnet 50 networks in Austin that use a 128 Kbps line, name the network type Austin.

By default, the management console assigns a client network to the Unassigned network type. If you are planning to implement network types, assign a network type to a client network that represents the common latency characteristics for a group of networks rather than the default network type.

Note: The easiest way to create client networks is to create an Excel chart of networks and associated network type information and import the list into the management console as a CSV file.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Network Types in the Show Me menu.
3. Under the Show Me menu, click Add Network Type.
4. Complete the fields in Network Type Properties and click OK.

For information about specifying network type properties, click Help.

More information:


[Import a CSV File](#) (see page 42)

[Assign a Network Type to a Client Network](#) (see page 54)

Edit a Network Type

Edit a network type to change its name or network incident response. Changes to a network type apply to corresponding client networks where the network type is assigned.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Network Types in the Show Me menu.
3. In Network Types, click the edit icon () to edit a network type.
4. Complete the fields in Network Type Properties and click OK.

For information about network type properties, click Help.


More information:

[Managing Incident Responses](#) (see page 165)

Delete a Network Type

If you delete a network type that is assigned to a network, the management console automatically reassigns the default network type, Unassigned, to the client network. Note that you can assign an incident response to the Unassigned Network Type.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Network Types in the Show Me menu.
3. Click the red X () to delete a network type.
4. Click Continue with Delete to confirm that you want to delete the network type.

The management console assigns the Unassigned network type to networks that were previously assigned to the deleted network type.

More information:

[Edit a Network Type](#) (see page 53)

Assign a Network Type to a Client Network

Assign a network type to the client networks where you want to launch a particular action in response to a network incident. We recommend assigning a network type to:

- A client network. When assigning an incident response to a network type, keep in mind that the management console initiates the specified actions in response to a network incident on any client networks where you assign the network type.

When you add a client network to the management console, be sure to also assign a network type. By default, the management console does not assign a network type to a new client network. If necessary, you can edit one or more client networks to assign a network type.

- A server subnet. The management console automatically assigns the specified network type to the /32 and /128 client networks that it creates from corresponding server traffic.

More information:

[How /32 Client Networks Work](#) (see page 70)

[Edit a Client Network](#) (see page 46)

[Edit a Server Subnet](#) (see page 74)

Manage Client Networks with Web Service Methods

Of the three primary configuration and reporting elements in the management console--applications, networks, and servers--networks are the most numerous and difficult to maintain.

Use Web service methods to create, read, update, and delete network definitions programmatically using the Network Configuration Service application programming interface (API) instead of adding networks through the management console.

See the following sections for more information about configuring networks using Web service methods.

Parameter Descriptions

The following sections discuss the parameters that are used throughout the described web service methods:

ClientId

Type: Unsigned 4-byte integer

Definition: The network definition identifier (clients table entry). Possibly mapped to more than one network (client_cache table entry), depending on the number of sub regions configured.

In the Administration page, each network definition in the Network List has a unique client ID, but it is not displayed.

ClientSetId

Type: Unsigned 4-byte integer

Definition: The network set to which this definition belongs. This field is not exposed and always defaults to the first active client set identifier.

Description

Type: String of up to 50 characters (latin1)

Definition: User-specified label for the network definition being added.

In the Administration page, the Network column in the Network List lists the description for each network definition.

NetworkType

Type: String of up to 50 characters (latin1)

Definition: Identifies the assigned type of the definition or empty/null/*Unassigned* if none is required. This parameter is case insensitive on matches. If a match is not found, a new network type will be added.

In the Administration page, the Network Type column in the Network List shows the network type associated with each network definition. If a definition is not assigned a network type, it is marked Unassigned.

Regions

Type: Integer power of 2 between 1 and 256. Regions is constrained by subnet mask selection; for example, /31 can only expand into 2 sub regions.

Definition: The number of sub regions into which this definition should be expanded.

In the Administration page, the Regions column in the Network List shows the number of sub regions into which each network definition will expand. Those displayed are set to the minimum; that is, one region per definition.

Subnet

Type: String in x.x.x.x/m format, where x.x.x.x is a valid IP address and m is an integer subnet mask between 1 and 32.

Definition: The address and mask of the subnet for this network definition.

In the Network List, the Subnet column shows the subnet of the network definition in the expected address/mask input and output format.

Web Service Methods

This section describes web service methods.

More information:

[How Error Reporting Works](#) (see page 61)

InsertNetworkDefinition

Purpose: Lets you create a network definition.

In the Administration page, each submitted, new entry in the Network Properties would be equivalent to one method call.

Input parameters:

- Description (optional)
- Subnet (required)
- Regions (required)
- NetworkType (optional, empty value equals *unassigned*)

Return: XML document with root node *InsertNetworkDefinition* and *ClientId* element containing the newly created ClientId for the inserted network. If unsuccessful, ClientId will be zero.

Errors returned: Standard error reporting is provided.

UpdateNetworkDefinition

Purpose: Lets you modify a network definition identified by the passed client ID, including the subnet description, IP address, and mask.

In the Administration page, each network definition edit is equivalent to one method call.

Input parameters:

- ClientId (required)
- Description (optional, but empty value erases old description)
- Regions (required)
- NetworkType (optional; empty value equals *unassigned*, new values create new NetworkType)

Return: True or false indicating success or failure of the update.

Errors returned: Standard error reporting is provided.

DeleteNetworkDefinition

Purpose: Lets you delete a network definition identified by the passed client ID. Internally, the client network is marked *inactive*. Deleting already deleted networks will result in a *false* return code.

In the Administration page, deleting a network from the Network List is equivalent to one method call.

Input parameter: ClientId (required)

Return: XML document with root node *DeleteNetworkDefinition* and *ClientId* element containing the deleted ClientId.

Errors returned: Standard error reporting is provided.

NetworkDefinitionsList

Purpose: Lets you retrieve network definitions by client ID.

In the Administration page, the Network List includes each network definition from the Default Network Set.

Input parameter: None.


Return: XML document with one entry per network definition. Each entry contains the following parameters:

- ClientId
- Description
- NetworkType
- Subnet
- Regions

Errors returned: Standard error reporting is provided.

ReloadCollectors

Purpose: Triggers a reset of all monitoring devices to pick up the configuration data and then restart.

In the Show Me menu of the Administration page, click Data Monitoring and Monitoring Devices. In the ADA Monitoring Device List, click the blue gear menu () and select Synchronize Monitor Devices.

Input parameter: None

Return: XML document with root node *ReloadCollectors* and *Collector* element for status report for each monitoring device. The Status/Message attribute indicates the success of the reload process. Check each monitoring device individually to find the results of the synchronization.

Errors returned: Standard error reporting is provided.

ShowVersion

Purpose: Shows the string/numerical version number of this API. This method is applicable only to this web services API.

Input parameter: None.

Return: Version number as a string.

Errors returned: Standard error reporting is provided.

How to Test the Web Services API

Test and access the Network Configuration Service API and manage your network definitions programmatically.

Follow these steps:

1. Navigate to <http://localhost/SuperAgentWebService/NetworkConfigService.asmx>. This site provides a convenient user interface into the NetworkConfigService and its Service Description.

To access the NetworkConfigService test site from an accessible remote computer, replace localhost with <ADA_Server_FQDN>.

2. Look in Application EventLog under SuperAgent NwkConfig WS for additional explanatory messages for failures.
3. For scripted operations, review the [sample perl script](#) (see page 62) to exercise the API.

More information:

[Sample Perl Script](#) (see page 62)

How Error Reporting Works

All interfaces except ShowVersion() return an XmlDocument object. The root node of the document is named after the method; for example, NetworkDefinitionsList or UpdateNetworkDefinition. The following attributes are attached to this root node:

- **Status:** A True or False string literal that indicates whether the execution was successful. For ReloadCollectors(), a True status does not necessarily mean the operation was successful because internally a reload might report no exception when the monitor is offline or unavailable.
- **Message:** A string containing the specific error message for a failure. Message will be empty if the operation is successful. The complete details with stack trace appears in the Windows EventLog.

The following examples illustrate the object and attributes:

```
<NetworkDefinitionsList Status="True" Message="">
  <Network>
    <ClientId>3</ClientId>
    <Description>192.168.0.2</Description>
    <SubnetMask>192.168.0.2/32</SubnetMask>
    <Regions>1</Regions>
  </Network>
</NetworkDefinitionsList>
```

```
<?xml version="1.0" encoding="utf-8" ?> <ReloadCollectors Status="False" Message="Can't connect to MySQL
server on 'localhost' (10061)" />
<?xml version="1.0" encoding="utf-8" ?> <UpdateNetworkDefinition Status="True"
Message=""><ClientId>7</ClientId></UpdateNetworkDefinition>
<?xml version="1.0" encoding="utf-8" ?> <InsertNetworkDefinition Status="True"
Message=""><ClientId>8</ClientId></InsertNetworkDefinition>
```

Sample Perl Script

```
#####  
#####  
#####  
# SAMPLE perl script exercising the Network Configuration API  
# CA  
#####  
#  
#####  
# Note: please modify $url to point to where ADA Console is hosted  
#####  
#  
#####  
#!/usr/local/bin/perl  
use Data::Dump qw(dump);  
use SOAP::Lite (  
#       +trace => all,  
        maptype => {}  
);  
$SOAP::Constants::DO_NOT_USE_CHARSET = 1;  
#  
#  
#####  
my $uri = "http://www.netqos.com/networkconfig";  
my $url = "http://localhost/SuperAgentWebService/NetworkConfigService.asmx";  
#  
my ($method, $result, $networks, @nodes, @params);  
my ($clientId, $description, $subnetCidr, $regionCount, $networkType);  
#  
sub SOAP::Transport::HTTP::Client::get_basic_credentials {  
    return $user => $pass;  
}  
#  
#####  
my $soap = SOAP::Lite  
    -> uri($uri)  
    -> on_action( sub { join ' ', @_ } )  
    -> proxy($url);  
#  
#  
#####  
# NetworkDefinitionsList  
#####  
$method = SOAP::Data->name('NetworkDefinitionsList')->attr({xmlns => $uri});  
$networks = $soap->call($method);  
print "\nNetworkDefinitionsList:\n";
```

```

if ($networks->fault) {
    print join ', ', $result->faultcode, $result->faultstring, $result->faultdetail;
} else {
    if ($networks->dataof('//NetworkDefinitionsList')->{_attr}->{Status} eq "False") {
        print $network->dataof('//NetworkDefinitionsList')->{_attr}->{Message}, "\n";
    } else {
        print "\n\nClientID\tDescription\tSubnet\t\tRegions\t\tNetworkType\n";
        @nodes = $networks->valueof('//Network');
        foreach $n (@nodes)
        {
            print $n->{ClientID}, "\t", $n->{Description}, "\t\t",
                $n->{Subnet}, "\t\t", $n->{Regions}, $n->{NetworkType}, "\t", "\n";
        }
    }
}
#
#####
# UpdateNetworkDefinition
#####
print "\n\nUpdateNetworkDefinition:\n";
$method = SOAP::Data->name('UpdateNetworkDefinition')->attr({xmlns => $uri});
$clientID = $networks->valueof('//Network/ClientID');
$description = $networks->valueof('//Network/Description') . " UPDATED";
$description = substr($description, 0, 50);
$regionCount = $networks->valueof('//Network/Regions');
$networkType = "unassigned";
@params = ( SOAP::Data->name(ClientID => $clientID),
            SOAP::Data->name(Description => $description),
            SOAP::Data->name(Regions => $regionCount),
            SOAP::Data->name(NetworkType => $networkType)
);
$result = $soap->call($method => @params);
if ($result->fault) {
    print join ', ', $result->faultcode, $result->faultstring, $result->faultdetail;
} else {
    if ($result->dataof('//UpdateNetworkDefinition')->{_attr}->{Status} eq "False") {
        print $result->dataof('//UpdateNetworkDefinition')->{_attr}->{Message};
    } else {
        print "UpdateNetworkDefinition($clientID, $description, $regionCount, $networkType):\n";
        print dump($result->result), "\n";
    }
}
#
#####
# DeleteNetworkDefinition
#####
print "\n\nDeleteNetworkDefinition:\n";
$method = SOAP::Data->name('DeleteNetworkDefinition')->attr({xmlns => $uri});
$clientID = $networks->valueof('//Network/ClientID');

```

```
$description = $networks->valueof('//Network/Description');
@params = (SOAP::Data->name(ClientId => $clientId)->attr({xmlns => $uri}));
$result = $soap->call($method => @params);
if ($result->fault) {
    print join ', ', $result->faultcode, $result->faultstring, $result->faultdetail;
} else {
    if ($result->dataof('//DeleteNetworkDefinition')->{_attr}->{Status} eq "False") {
        print $result->dataof('//DeleteNetworkDefinition')->{_attr}->{Message};
    } else {
        print "DeleteNetworkDefinition($clientId, $description):\n";
        print dump($result->result), "\n";
    }
}
#
#####
# InsertNetworkDefinition
#####
print "\n\nInsertNetworkDefinition:\n";
$method = SOAP::Data->name('InsertNetworkDefinition')->attr({xmlns => $uri});
$description = $networks->valueof('//Network/Description');
$subnetCidr = $networks->valueof('//Network/Subnet');
$regionCount = $networks->valueof('//Network/Regions');
$networkType = "unassigned";
@params = ( SOAP::Data->name(Description => $description),
            SOAP::Data->name(Subnet => $subnetCidr),
            SOAP::Data->name(Regions => $regionCount),
            SOAP::Data->name(NetworkType => $networkType)
);
$result = $soap->call($method => @params);
if ($result->fault) {
    print join ', ', $result->faultcode, $result->faultstring, $result->faultdetail;
} else {
    if ($result->dataof('//InsertNetworkDefinition')->{_attr}->{Status} eq "False") {
        print $result->dataof('//InsertNetworkDefinition')->{_attr}->{Message};
    } else {
        print "InsertNetworkDefinition($description, $subnetCidr, $regionCount, $networkType):\n";
        print dump($result->result), "\n";
    }
}
#
#
#####
# ReloadCollectors
#####
$method = SOAP::Data->name('ReloadCollectors')->attr({xmlns => $uri});
print "\n\nReloadCollectors():...\n";
$result = $soap->call($method);
if ($result->fault) {
    print join ', ', $result->faultcode, $result->faultstring, $result->faultdetail;
```



```

} else {
  if ($result->dataof('//ReloadCollectors')->{_attr}->{Status} eq "False") {
    print $result->dataof('//ReloadCollectors')->{_attr}->{Message};
  } else {
    @nodes = $result->valueof('//Collector');
    foreach $node (@nodes)
    {
      print $node->{Address}, "\t", $node->{Status}, "\t", $node->{Info}, "\n";
    }
  }
}
#
#
#####
# ShowVersion
#####
$method = SOAP::Data->name('ShowVersion')->attr({xmlns => $uri});
print "\n\nShowVersion():\n";
$result = $soap->call($method);
if ($result->fault) {
  print join ', ', $result->faultcode, $result->faultstring, $result->faultdetail;
} else {
  print $result->result, "\n";
}
#

```


Chapter 3: Managing Servers

This section contains the following topics:

[How Servers Work](#) (see page 67)

[Manage Server Subnets](#) (see page 72)

[Manage Servers](#) (see page 76)

[Pinning a Monitor Feed to a Server](#) (see page 86)

[Schedule Server Maintenance](#) (see page 87)

How Servers Work

A *server* specifies a server IPv4 address that you want to monitor. The management console works best when configured to monitor server subnets. A server subnet specifies a range of server IPv4 addresses that you want the management console to monitor, and enables the management console to automatically monitor application traffic on matching servers.

By default, the management console automatically monitors a pre-defined server subnet called *Monitor All Servers* on subnet 0.0.0.0 mask 0. The default *Monitor All Servers* subnet enables automatic data collection when adding a new monitor.

More information:

[How the Management Console Manages Database Growth](#) (see page 246)

How the Server Subnet List Works

Use the Server Subnet List to specify a server subnet which identifies a range of server IP addresses to be monitored by the management console. Based on the server subnets you specify, when a monitoring device observes matching application traffic between the server and a client network, the management console assigns the monitoring device to the server and adds the server to the Server List. You can also manually add a server to the Server List.

By default, the management console automatically monitors a pre-defined server subnet called *Monitor All Servers* on subnet 0.0.0.0 mask 0. The default *Monitor All Servers* subnet enables automatic data collection when adding a new monitor.

[Go to Server List](#)

Server Subnet List

View and manage server subnets to specify the range of server IP addresses monitored by SuperAgent. When matching traffic on a server is observed, the server is added to the Server List.

Description	Subnet	Exclusions	Apps	Servers		
CA Client Network	130.200.39.0/24	0	1	1	<input type="checkbox"/>	
My_subnet	10.0.38.0/24	0	0	0	<input type="checkbox"/>	
NQ Client Network	192.168.0.0/24	0	13	7	<input type="checkbox"/>	

[Go to Server Subnet List](#)

Server List

View and manage the servers that are monitored by SuperAgent. If you do not see a particular server, go to the Server Subnet List and verify that a matching server subnet exists.

Server	Address	Collection Device	Apps	Bytes	User Modified	Last Seen			
130.200.39.244	130.200.39.244	Stndalone32-W01 - Primary SPAN Feed		0	0	Yes	Inactive	<input type="checkbox"/>	
192.168.0.2	192.168.0.2	Stndalone32-W01 - Primary SPAN Feed		0	0	No	Inactive	<input type="checkbox"/>	
192.168.0.20	192.168.0.20	Stndalone32-W01 - Primary SPAN Feed		1	1.8M	No	Current	<input type="checkbox"/>	

How the Server List Works

Use the Server List to view and manage the servers that are monitored by the management console. The management console automatically monitors all application traffic on each server in the Server List.

Typically, the Server List is populated with servers that correspond to a specified range of IP addresses in the Server Subnet list. If necessary, you can add a particular server that you want the management console to monitor -- the management console immediately adds the server to the Server List.

[Go to Server List](#)

Server Subnet List							
<i>View and manage server subnets to specify the range of server IP addresses monitored by SuperAgent. When matching traffic on a server is observed, the server is added to the Server List.</i>							
<input type="text" value="Add Server Subnet"/>							
Description	Subnet	Exclusions	Apps	Servers			
CA Client Network	130.200.39.0/24	0	1	1	<input type="checkbox"/>		
My_subnet	10.0.38.0/24	0	0	0	<input type="checkbox"/>		
NQ Client Network	192.168.0.0/24	0	13	7	<input type="checkbox"/>		

[Go to Server Subnet List](#)

Server List							
<i>View and manage the servers that are monitored by SuperAgent. If you do not see a particular server, go to the Server Subnet List and verify that a matching server subnet exists.</i>							
<input type="text" value="Add Server"/>		<input type="text" value="Search"/>		<input type="text" value="Clear"/>			
Server	Address	Collection Device	Apps	Bytes	User Modified	Last Seen	
130.200.39.244	130.200.39.244	Stndalone32-W01 - Primary SPAN Feed		0	0	Yes	Inactive <input type="checkbox"/>
192.168.0.2	192.168.0.2	Stndalone32-W01 - Primary SPAN Feed		0	0	No	Inactive <input type="checkbox"/>
192.168.0.20	192.168.0.20	Stndalone32-W01 - Primary SPAN Feed		1	1.8M	No	Current <input type="checkbox"/>

When registered as a data source with the CA PC or the CA NPC, the CA PC or CA NPC automatically groups the servers for reporting.

More information:

[Manage Server Subnets](#) (see page 72)

[System Groups](#) (see page 215)

How /32 Client Networks Work

When a server is added to the Server List, the management console automatically creates a corresponding /32 client network. A /32 client network is a client network with a single IPv4 address.

Use /32 client networks to monitor a multi-tier application. A multi-tier application is an application with more than one server, and communication between servers is performed by at least one server that acts as both a server to client requests, and a client to another server.

To enable a management console user to quickly analyze and respond to network issues on a multi-tier application, assign a network type to your /32 client networks.

When you add a server subnet, you can assign a default network type so that when the management console creates a corresponding /32 client network, the correct network type is assigned.

You cannot delete a /32 client network that was created by the management console. The management console automatically deletes a /32 client network when its corresponding server is deleted from the Server List.

More information:

[How Client Networks Work](#) (see page 29)

[Manage a Multi-Tiered Application](#) (see page 130)

TCP Session Identification

To enable the management console to accurately report application response times from the server, a monitoring device must be able to identify the server-side of a client/server TCP session. A monitoring device uses the following criteria to determine which of the two endpoints in a TCP session represents the server side:

- New TCP session. To identify the server-side in a new TCP conversation, the IP address that is the SYN recipient is checked against the server subnets definition. The IP address of the SYN recipient must be in one of the specified server subnets. Both the SYN and SYN-ACK must be observed.
- Existing TCP session (traffic for conversations where the TCP connection setup was not seen). If neither endpoint IP falls within the included range of one of the server subnets, then the conversation is ignored.

If only one of the two IP addresses lies within the range of a server subnet, then that endpoint is presumed to be the server.

If both IP addresses lie within the range of one or more server subnets and neither is recorded as a well-known application, the monitoring device assumes that the lowest port is the server.

Host Name Resolution

If the management console is registered as a data source with the CA PC or the CA NPC, the CA PC or CA NPC uses a proxy server to automatically query DNS when the management console adds a server to the Server List. Alternatively, the management console sends lookup requests to its DNS server on the default port, UDP-53.

The management console does not automatically resolve the host name of a server when you:

- Add a server manually
- Import a list of servers

To manually resolve the host name of a server, edit the server properties.

More information:

[Manage Servers](#) (see page 76)

[Manage Console Settings](#) (see page 223)

Manage Server Subnets

A server subnet specifies a contiguous range of server IP addresses that you want to monitor. When a monitoring device observes application traffic that matches a specified server subnet, the management console:

- Adds the server to any matching System applications.
- Adds the server to the Server List, where you can view and manage the properties of the server, for example, to assign an incident response.
- Adds the server to the Network List as a /32 or /128 client network.

We recommend adding server subnets to CA Application Delivery Analysis that closely align with your server VLAN definitions. To limit the TCP ports that the management console monitors on each server subnet, assign port exclusions to the server subnet.

Application Delivery Analysis includes a default server subnet called *Monitor All Servers* on subnet 0.0.0.0 mask 0. The default server subnet enables automatic data collection when adding a new monitor.

More information:

[How /32 Client Networks Work](#) (see page 70)

[Application Port Exclusions](#) (see page 106)

[How Applications Work](#) (see page 103)

[How the Management Console Manages Database Growth](#) (see page 246)

[Manage Servers](#) (see page 76)

Add a Server Subnet

Add a server subnet to automatically monitor matching server traffic. When specifying a server subnet, specify a range of IP addresses that represents the application servers that you want the management console to monitor. Typically, this range closely aligns with your server VLAN definitions.

After you add a server subnet, edit the server subnet to add an exclusion to ignore traffic from particular servers within the specified range of server IP addresses.

To specify servers you want to	Specify a subnet mask between
Include	■ IPv4: /16 and /31. You cannot specify a /32 subnet mask.
Exclude	■ IPv4: /17 and /32

If necessary, you can add a server to the management console before a monitoring device observes its traffic.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Servers in the Show Me menu.
3. (Optional) Choose the domain where you want to add the server subnet.
4. Scroll to the Server Subnet List and click to add an IPv4 server subnet.

Add Server Subnets opens.

5. Complete the fields in Server Subnets and click OK.

For information about setting server subnet properties, click Help.

The management console validates the server subnet does not conflict with any existing server subnets you have defined and then adds the server subnet.

6. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

After 5-10 minutes, the Server List automatically displays servers that match the specified server subnet.

More information:

[Managing Tenants](#) (see page 95)


[Add a Server](#) (see page 78)

Edit a Server Subnet

Edit a server subnet to change its range of IP addresses, for example, to exclude particular servers, or to specify a default network type for the client networks that the management console automatically creates from matching server traffic.

Existing data that is outside an updated IP range continues to be available for reporting purposes.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Servers in the Show Me menu.
3. (Optional) Choose the domain with the server subnet you want to edit.
4. Scroll to the Server Subnet List and click  to edit a server subnet.

The Server Subnet Properties opens.

5. Complete the fields in Server Subnet Properties and click Apply.

For information about setting server subnet properties, click Help.

The management console validates that the server subnet does not conflict with any existing server subnets you have defined and then adds the server subnet.

6. Click Add Exclusion to exclude particular IP addresses within the range of IP addresses specified by the server subnet.

Specify a /16 (or higher) subnet mask.

7. Click OK.
8. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

More information:

[Managing Tenants](#) (see page 95)


[How /32 Client Networks Work](#) (see page 70)

Delete a Server Subnet

Delete a server subnet to prevent the management console from automatically monitoring new servers that match the server subnet. The management console continues to monitor existing servers that match the server subnet.

If necessary, delete servers from the Server List. Existing data continues to be available for reporting purposes.

Follow these steps:

1. Click the Administration page
2. Click Data Monitoring, Servers in the Show Me menu.
The Server List appears.
3. (Optional) Choose the domain with the server subnet you want to edit.
4. Click  to delete a server subnet in the Server Subnets List.
5. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

More information:

[Managing Tenants](#) (see page 95)

Manage Servers

Use the Servers List to manage server properties, such as:

- **SNMP profile.** The management console requires a SNMP profile with valid SNMP user credentials to perform a Performance via SNMP investigation on a server.

If you do not specify a valid SNMP profile, the management console attempts to discover the SNMP profile.

- **Monitor feed assignment.** You can override the monitor feed assignment that is automatically maintained by the management console to assign a particular monitor feed to a server, and choose to permanently override the assignment.

If you want the management console to monitor server traffic from more than one monitor feed, for example when the network fails over to a different network path, or server traffic is load-balanced between two switches, you can assign a secondary monitor feed to the server's assigned monitor feed.

- **TCP/IP Host Name Resolution.** The management console automatically resolves the DNS host name for a server that matches a specified server subnet.

For manually added or imported servers, edit the server properties to enter the host name.

- **Server maintenance schedule.** The management console assigns a default maintenance schedule to a server, but you must assign a scheduled period for maintenance.

- **Incident response.** By default, the management console does not launch an action in response to a server incident.

If you plan to enable a Performance via SNMP investigation, it is recommended that you also assign an SNMP profile to the server.

To limit the TCP ports that are monitored by the management console, assign a server or a server subnet to a port exclusion.

More information:

[Add an Action to a Network or Server Incident Response](#) (see page 177)

[Create a Pair of Monitor Feeds](#) (see page 243)

[How Monitor Feed Assignment Works](#) (see page 241)

[Manage SNMP Profiles](#) (see page 225)

[Schedule Server Maintenance](#) (see page 87)


Naming Conventions

The following table lists suggested naming conventions for servers.

Server Type	Suggested Naming Convention	Example
Single function server	<i>DNSName-IPAddress</i>	Goliath-196.128.34.1
One application, multiple servers in a farm	<i>ApplicationName-DNSName</i>	DocMgr-Zeus DocMgr-Athena DocMgr-Mercury
One application, multiple servers, multiple locations	<i>ApplicationName-DNSName-Location</i>	DocMgr-Hamlet-NewYork DocMgr-Romeo-Milan DocMgr-Othello-London

Find a Server

Search the Server List to find servers that are currently monitored.

In the Server List, click the magnifying glass  icon to view the applications that the management console is monitoring on a server.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Servers in the Show Me menu.
3. (Optional) Choose the domain with the server subnet you want to edit.
4. Scroll to the Server List, enter a search string in the field, and click Search.
Pattern-matching characters such as * or % are not supported.
5. Click Clear to reset the list.

More information:

[Managing Tenants](#) (see page 95)

Add a Server

The management console automatically populates the Server List with observed servers that match a specified server subnet and have application port traffic that does not match any port exclusions you have defined.

If you want to assign a server to an application before the management console has observed its traffic, add the server and then assign the server to the application.

When adding a server, avoid overriding the server's automatic monitor feed assignment. Instead, allow the management console, if necessary, to automatically re-assign the monitor feed.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Servers in the Show Me menu.
3. (Optional) Choose the domain where you want to add the server.
4. Scroll to the Server List and click to add an IPv4 server.

Server Properties opens.

5. Complete the fields in Server Properties and click OK.

For information about setting server properties, click Help.

6. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.


Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

More information:


[Managing Tenants](#) (see page 95)


Edit a Server

Use the Server List to view and manage the servers that are monitored by the management console:

- Click the magnifying glass  icon in the Apps column to view the applications that the management console is monitoring on a particular server.
- Edit a server, for example, to view its monitor feed statistics or change its properties. If you change the properties of a server, its User Modified status changes to Yes. Monitor feed statistics include:
 - Server Traffic by Feed (Last 24 Hours) Describes the traffic volume by a server on each monitor feed, which enables you to verify that the monitor feed assignment is correct for both automatically assigned and pinned monitor assignments.
 - Volume Statistics for Assigned Monitor Device (Last 7 Days) Describes the traffic volume observed by a server's assigned monitoring device from the past 7 days.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Servers in the Show Me menu.
3. (Optional) Choose the domain with the server you want.
4. Scroll to the Server List and click  to edit a server.

(Optional) To edit more than one server, select each server and then click  to bulk edit all the selected servers. Note that any changes you make apply to all selected servers.

The Server Properties opens.

5. Complete the fields in Server Properties and click OK.
For information about setting server properties, click Help.
6. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

More information:

[Managing Tenants](#) (see page 95)


Delete a Server

Delete a server to stop monitoring its application port traffic. After you delete a server, existing data continues to be available for reporting purposes.

Deleting a server that matches a server subnet will only temporarily delete the server. When a monitoring device observes server traffic that matches the server subnet, the management console will subsequently add the server back to the list. If the server you want to delete matches a server subnet, edit the server subnet and add a server exclusion to prevent the management console from monitoring the server.

If you do not want a particular port or range of ports on the server, add a port exclusion.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Servers in the Show Me menu.
3. (Optional) Choose the domain with the server you want to delete.
4. Scroll to the Server List and click  to delete a server.
5. Click Continue with Delete at the prompt to delete the server.
6. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

More information:

[Managing Tenants](#) (see page 95)

[Application Port Exclusions](#) (see page 106)

[Edit a Server Subnet](#) (see page 74)

Import Server Definitions from a CSV File

To quickly add new servers to the management console, add a comma-separated list of server host names and IPv4 addresses to a CSV file, and then import the file into the management console. You cannot:

- Import a server subnet definition
- Re-import an existing IPv4 address

As part of the import process, you can bulk edit the list of servers, for example, to assign a server maintenance schedule. When you are finished, synchronize the monitoring devices to begin monitoring the servers using the properties you specify.

Create a CSV File

To get started quickly, add a comma-separated list of server names and IPv4 addresses to a CSV file.

If you do not include both a server name and IP address, you can edit the list of definitions before you complete the import. For example, you can create a CSV file with IP addresses only, and then specify the corresponding server names before you complete the import. The management console does not query DNS to resolve host names for imported servers.

Follow the steps below to create a CSV file that specifies the servers that you want the management console to monitor. When creating a CSV file, separate each field with a comma and enclose:

- Embedded commas or double quotes inside double quotes.
- Strings that contain one or more spaces in double quotes; for example: "Houston Office".

An example of CSV-formatted server definitions appears below. The Austin DNS Server entry specifies that ref-sa-coll, which is a CA Standard Monitor, sees the server traffic for 192.168.100.2 on its Packets monitor feed and, by default, should use the SNMP profile named "netqos" to poll the server:

"Austin DNS Server", 192.168.100.2, "ref-sa-coll", "Packets", "netqos"

Follow these steps:

1. Create a text file with a .CSV extension, for example, austin_servers.csv.
2. Specify each comma-delimited server definition on a separate line.
3. Use the following format for each entry:

Server_Name,Server_IP,Monitoring_Device,Monitor_Feed,SNMP_Profile

Server_Name

Defines the name of the server, up to 50 characters, as you want it to appear in the management console.

Note that when you import a list of servers, the management console does not automatically resolve the DNS host name.

Server_IP

Defines the IPv4 address of the server in four-part, dotted-decimal notation. If you do not specify the address format correctly, the import uses the IP address as the server name.

Monitoring_Device

(Optional) Defines the monitoring device that sees the server's traffic. If you specify the monitoring device, you must also specify a monitor feed on the device. After you import the server, if necessary, the management console automatically reassigns the monitoring device and monitor feed. If you do not specify a monitor feed, the management console automatically assigns the best monitor feed.

Monitor_Feed

(Optional) Defines the monitor feed on the monitoring device that sees the server's traffic. If you specify the monitoring device, you must also specify a monitor feed on the device, such as Packets for the Packets monitor feed. After you import the server, if necessary, the management console automatically reassigns the monitoring device and feed.

If you do not specify a monitor feed, the management console automatically assigns the best monitor feed.

To browse the list of available monitor feed names on a monitoring device, edit the CA Standard Monitor or CA Multi-Port Monitor.

SNMP_Profile

(Optional). Defines the SNMP profile name that you want the management console to use when polling the server, for example, as part of a Performance via SNMP investigation. Specify a null value ("") if you want the management console to discover a valid SNMP profile.

4. Save the file and import it into the management console.

More information:


[How Monitoring Devices Work](#) (see page 239)

[How SNMP Profile Discovery Works](#) (see page 226)

Import Server Definitions

When importing server definitions, make sure the text file has a .CSV file extension. As part of the import, you can bulk edit the server definitions, for example, to assign a SNMP profile.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Servers in the Show Me menu.
3. (Optional) Choose the domain into which you want to import servers.
4. Scroll to the Server List and click the blue gear menu () to select Import from File.

Import Server Definitions opens.

5. Review the information about the CSV file format and if necessary, copy and paste the example into your CSV file.
6. Click Browse to select the CSV file, and click Next.

Save Import Server Definitions opens.

7. Review the list of imported server definitions and, if necessary, make any changes, and click OK.
8. Resolve any issues with the definitions you imported from the CSV file and the existing server definitions in the management console.


Save Imported Server Definitions highlights any conflicting server definitions.

9. Edit the highlighted server definitions to resolve any issues and click OK.

Yellow

Indicates the specified monitoring device does not exist or the server has an invalid IP address. The management console lets you proceed with the import, but before you proceed, you should resolve these types of issues.

Red

Indicates the server IP address is a duplicate. To proceed with the import, click  to delete any duplicate IP addresses.

10. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.


More information:

[Managing Tenants](#) (see page 95)

Export Server Definitions to a CSV File

Export IPv4 server definitions to a .CSV file for use as a template for importing new server definitions. Note that you cannot re-import an existing IPv4 server address.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Servers in the Show Me menu.
3. (Optional) Choose the domain from which you want to export servers.
4. Scroll to the Server List and click the blue gear menu () to select Export Servers.

Export Server Definitions opens.

5. Choose an export option and click OK.

Only unique IP addresses

Specifies that you want to export the IP address only.

All export fields

Specifies that you want to export the server host name and IP address, plus the assigned monitoring device and monitor feed, and the assigned SNMP profile name.

The File Download dialog box opens.

6. Click Open to view the CSV file contents or click Save to save the CSV file.

More information:

[Managing Tenants](#) (see page 95)

Pinning a Monitor Feed to a Server

Pin a monitor feed to permanently assign a particular monitor feed to a server, for example, when you know the monitoring device from which you want the management console to monitor the server, and you do not want the management console to change it.

When you delete a monitoring device, any servers that were pinned to the corresponding monitor feed are unpinned, and another monitor feed is automatically assigned. It can take up to 10 minutes to update the monitor feed assignment.

When you add a server, you can choose to pin a monitor feed. Or, edit the server properties to assign a particular monitor feed to a server.

More information:

[Edit a Server](#) (see page 79)

Schedule Server Maintenance

A server maintenance schedule specifies when server performance is expected to be abnormal, for example, because of scheduled maintenance tasks, such as backups or software updates.

During a period of scheduled server maintenance, the management console:

- Does:
 - Close existing server incidents when the maintenance period begins.
If the server incident condition continues to exist after the maintenance period ends, the management console opens a new server incident.
 - Continue to collect data on application, server, and network performance, and rates performance as Normal, Minor (yellow), or Major (orange). This data can be useful for understanding performance before, during, and after a maintenance period.
- Does not:
 - Trigger a response to a new server incident created during a maintenance period.
 - Calculate sensitivity thresholds with data collected during a maintenance period.
 - Calculate baselines with data collected during a maintenance period.
 - Calculate performance OLAs or availability OLAs.




Note: A management console user can choose whether to report on degraded application, server, and network performance during a maintenance period. In the report Settings, use the Include Scheduled Maintenance option to show or hide performance data collected during a maintenance period.

How Maintenance Schedules Work

Use the Maintenance Schedules list to:

- Determine whether a maintenance schedule has at least one maintenance period
- View the number of servers to which a maintenance schedule is assigned

The management console provides a Default and Weekends maintenance schedule, however, these maintenance schedules do not include any maintenance periods. In the following example, the management console administrator has assigned at least one maintenance period to each of the maintenance schedules.

Maintenance Schedules					
Name	Periods	Servers	Edit	Delete	
Default	2	19			
Weekends	1	0			

The management console automatically assigns the Default maintenance schedule to newly observed servers. We recommend that you add a maintenance period to the Default maintenance schedule and, if necessary, assign a custom schedule to a server.

To view the maintenance schedule that is assigned to a particular server, edit the server properties.

More information:

[Edit a Server](#) (see page 79)

[Add a Maintenance Period to a Maintenance Schedule](#) (see page 91)

Add a Maintenance Schedule

Override the default maintenance schedule by adding a maintenance schedule with the appropriate maintenance periods, and assign the maintenance schedule to the appropriate servers.

Important! Make sure you define a maintenance period for the default maintenance schedule. The management console automatically assigns the default maintenance schedule to new servers.

After you add a maintenance schedule, assign a maintenance period to the schedule.

Follow these steps:

1. Click the Administration page.
2. Click Policies, Maintenance Schedules in the Show Me menu.
3. Under the Show Me menu, click Add Maintenance Schedule.

The Schedule Properties opens.

4. Type a schedule name and click Apply.

You are now ready to add a maintenance period to the maintenance schedule.

5. When you finish, click OK.

More information:


[Add a Maintenance Period to a Maintenance Schedule](#) (see page 91)

Rename a Server Maintenance Schedule

Rename a server maintenance schedule to a name that corresponds to your existing server maintenance schedules. Note that you can rename the Default and Weekends maintenance schedules that are provided with the management console.

While you are renaming a maintenance schedule, you can also modify its maintenance periods, for example, to add a maintenance period.

Follow these steps:

1. Click the Administration page.
2. Click Policies, Maintenance Schedules in the Show Me menu.
Maintenance Schedules opens.
3. Click  to edit a maintenance schedule.
4. Click Edit Schedule Name in the third Show Me menu to change the name of the maintenance schedule.
5. Type a new name in the Schedule Properties and click OK.

In Maintenance Schedules, the renamed maintenance schedule is displayed.

More information:

[Add a Maintenance Period to a Maintenance Schedule](#) (see page 91)

[Delete a Maintenance Period](#) (see page 92)


[Edit a Maintenance Period](#) (see page 92)

Delete a Maintenance Schedule

If you delete a maintenance schedule that is in use, the management console automatically assigns the Default maintenance schedule to any affected servers. To determine whether a maintenance schedule is assigned to a server, view the server properties.

If you have not done so already, it is recommended that you add a maintenance period to the Default maintenance schedule.

Follow these steps:

1. Click the Administration page.
2. Click Policies, Maintenance Schedules in the Show Me menu.
3. Click  to delete a maintenance schedule in Maintenance Schedules.
4. Click Continue with Delete at the prompt to delete the maintenance schedule.

More information:

[Edit a Server](#) (see page 79)

[Add a Maintenance Period to a Maintenance Schedule](#) (see page 91)





Add a Maintenance Period to a Maintenance Schedule

Add one or more maintenance periods per day to a maintenance schedule. After you add a maintenance period to a maintenance schedule, you are ready to assign the maintenance schedule to a server.

A maintenance period cannot cross the midnight boundary. For example, if you have a weekly maintenance period from 10 p.m. Saturday until 4 a.m. Sunday, you must create two maintenance periods:

- 10 p.m. to midnight Saturday
- Midnight to 4 a.m. Sunday


We recommend that you assign at least one maintenance period to a maintenance schedule. In the following example, the warning icon in the Periods column indicates that a maintenance period is not defined for the default maintenance schedule.

Maintenance Schedules					
Name	Periods	Servers	Edit	Delete	
Default	 0	21			
Mon from 22:00 to 23:00 CST6CDT	1	1			

Follow these steps:

1. Click the Administration page.
2. Click Policies, Maintenance Schedules in the Show Me menu.

The Maintenance Schedules list opens.

3. Click  to edit a maintenance schedule.

The Scheduled Periods list opens.

4. Click Add Period under the Show Me menu.

Scheduled Period Properties opens.

5. Specify the schedule period settings and click OK.

For information about specifying schedule period settings, click Help.

More information:

[Assign a Maintenance Schedule to a Server](#) (see page 93)


Edit a Maintenance Period

Edit a maintenance period to specify when server maintenance is planned to occur.


Follow these steps:

1. Click the Administration page.
2. Click Policies, Maintenance Schedules in the Show Me menu.

Maintenance Schedules opens.

3. Click  to edit a maintenance schedule.

Scheduled Periods opens.

4. Click  to edit a maintenance period.

Schedule Period Properties opens.


5. Specify the schedule period settings and click OK.

For information about specifying schedule period settings, click Help.


Delete a Maintenance Period

Delete a maintenance period to remove it from a server maintenance schedule. If you have assigned a maintenance schedule to a server, the schedule should include at least one valid maintenance period, however, the management console does not require an assigned maintenance schedule to have a maintenance period.

Follow these steps:

1. Click the Administration page.
2. Click Policies, Maintenance Schedules in the Show Me menu.
3. Click  to edit a maintenance schedule in the list of Maintenance Schedules.

The list of Scheduled Periods opens.

4. Click  to delete a maintenance period.

5. In the Delete Confirmation, click Continue with Delete to delete the maintenance period.

In Scheduled Periods, the maintenance period is removed.


More information:

[How Maintenance Schedules Work](#) (see page 88)


Assign a Maintenance Schedule to a Server

A server maintenance schedule identifies one or more maintenance periods where server performance is expected to be abnormal because of scheduled maintenance activities. Assign a maintenance schedule to one or more servers to specify when server maintenance is planned to occur.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Servers in the Show Me menu.
3. (Optional) Choose the domain with the server you want.
4. Scroll to the Server List and click  to edit a server.

Server Properties opens.

(Optional) To edit more than one server, select each server and then click  to bulk edit all the selected servers. Note that any changes you make apply to all selected servers.

5. Click Maintenance Schedule, choose a schedule from the list, and click OK.

For more information about server properties, click Help.

When editing more than one server, any changes you make apply to all servers. A value of No Change indicates that the existing values on each server will be preserved.

6. Click OK.
7. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

More information:

[Managing Tenants](#) (see page 95)

[Add a Maintenance Period to a Maintenance Schedule](#) (see page 91)


Chapter 4: Managing Tenants

Use domains to manage tenant data in CA Application Delivery Analysis. CA Application Delivery Analysis is not tenant-aware, however, you can separate tenant traffic by domain. In non-ISP environments, domains separate application traffic with overlapping (duplicate) client network IP addresses.

Important: In CA Application Delivery Analysis, the Administrator role has access to all domain data and domain configuration. Do not give a tenant user the Administrator role on CA Application Delivery Analysis.

Introduction to Tenancy

Use domains to uniquely identify traffic between a server and client network on a monitor feed. Separating tenant data by domain is how CA Application Delivery Analysis supports multi-tenancy.



Complete the following tasks:

1. [Add client networks to a domain](#) (see page 98).
2. [Add servers to a domain](#) (see page 99).
3. [Assign domains to monitor feeds](#) (see page 100).

Prerequisites

To manage tenants in CA Application Delivery Analysis with CA PC, make sure that:

- CA PC meets the following requirements:
 - CA Application Delivery Analysis is a registered data source.
 - Tenant users have permission to the appropriate IP domain groups. Do not give a tenant user access to domain data from other tenants.
 - Tenant users do not have the Administrator role on CA Application Delivery Analysis. The Administrator role gives a CA Application Delivery Analysis user access to all data across all domains and the Administration page.
 - The CA Application Delivery Analysis data source is synchronized.
- CA Application Delivery Analysis meets the following requirements:
 - The Monitor NIC on the monitoring device can read VLAN tag information from the packet header. If you do not need to separate traffic by VLAN, this is not applicable.
Note: Use the Configuration Utility to confirm that VLAN tag information is available, or, take a packet capture from the monitoring device.
 - The management console is synchronized with the list of domains in CA PC or the CA NPC.
Note: Click Data Monitoring, Domains on the Administration page to display a list of the available domains.

To manage tenants in CA Application Delivery Analysis with CA NPC, make sure that:

- CA NPC meets the following requirements:
 - CA Application Delivery Analysis is a registered data source.
 - Tenant users have permission to the appropriate domain groups. Do not give a tenant user access to domain data from other tenants.
 - Tenant users do not have the Administrator role on CA Application Delivery Analysis. The Administrator role gives a CA Application Delivery Analysis user access to all data across all domain groups and the Administration page.
 - The CA Application Delivery Analysis data source is synchronized.
- CA Application Delivery Analysis meets the following requirements:
 - The management console is synchronized with the list of domain groups in CA NPC.

How Domains Separate Traffic

Domains separate traffic at the:

- Monitor feed
- Client network
- Server subnet and server

With the same domain assigned to a server, network, and monitor feed, the management console reports the unique application traffic between a client and server from the monitor feed.

VLAN-tagged traffic can be separated at the monitor feed into different domains based on the VLAN tag definition. Separating VLAN-tagged traffic into domains enables a single monitor feed to monitor multiple domains.

Applications are domain-independent. Reports can show application performance across domains without requiring multiple application definitions, such as Exchange Company A and Exchange Company B.

Note that domains do not apply to the application properties. For example, if you rename an application, the application name is the same across domains. However, if you want to set different thresholds for application performance, performance OLAs, and availability OLAs, you must create an application for each domain.

Data Source Synchronization

CA PC and CA NPC synchronize a list of domains with CA Application Delivery Analysis. It can take up to 5 minutes for the management console to update its domain list.

When a domain is deleted from the CA PC, the management console:

- Deletes all client networks, server subnets, servers, and port exclusions associated with the domain, and removes associated server assignments from user-defined applications.

Existing domain-specific data continues to be available for reporting purposes.

Note that a user-defined application is not associated with a particular domain, so after you delete a domain, the management console continues to report on the application, if the management console observes application traffic between the servers and client networks that belong to any remaining domains.

- Reassigns all monitor feeds associated with the domain to the default domain. Note that you cannot delete the default domain.

More information:

[Users and Groups](#) (see page 214)

How Domain-Based Reporting Works

Use the Settings page to report on the servers and networks that correspond to a particular domain. Applications are domain-independent, therefore, you are not required to filter on an application by domain.

Use domain group permissions in CA PC and CA NPC to give a tenant user permission to the report data for their domain.

The report settings you choose are preserved when you drill-in to CA Multi-Port Monitor.

If the Show Me box in the management console does not display a Domains column, verify the prerequisites for managing tenants.

More information:

[Prerequisites](#) (see page 96)

Add Client Networks to a Domain

Add client networks to a domain. With the same domain assigned to a server, network, and monitor feed, the management console reports the unique application traffic between a client network and server.

When you add a client network, make sure to specify the correct domain. After you add a client network, you cannot change its domain. If necessary, delete the client network and then add it to the correct domain.

More information:

[Add a Client Network](#) (see page 44)

Add Servers to a Domain

Add servers to a domain. With the same domain assigned to a server, network, and monitor feed, the management console reports the unique application traffic between a client network and server.

When you add a server, make sure to specify the correct domain. After you add a server, you cannot change its domain. If necessary, delete the server and then add it to the correct domain.

More information:

[Add a Server Subnet](#) (see page 72)

[Add a Server](#) (see page 78)

Assign Domains to Monitor Feeds

Assign domains to specify where the management console reports:

- Untagged traffic. All untagged traffic is assigned to the domain for the monitor feed.
- VLAN-tagged traffic. Separate VLAN-tagged traffic on the monitor feed by assigning VLANs to domains.

You also can assign a domain to unassigned VLAN traffic.

By default, all traffic on the monitor feed is assigned to the Default Domain.

Important! When using a Multi-Port Monitor with CA CEM TIM, do not assign VLAN traffic on the logical port for the TIM to domains. The TIM does not support VLAN-based monitor feeds.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
3. Scroll to the ADA Monitoring Device List and click the edit icon (✎) to edit a monitoring device.

The Monitor Properties opens.

4. Scroll down to Monitor Feeds.
5. Click the edit icon (✎) to edit a monitor feed.

The Monitor Feed list expands.

- a. Click Domain to assign a domain to traffic that is not tagged with VLAN information.
 - b. Click Apply.
6. Click Assign VLANs.

In Assign VLANs, edit the VLAN assignments for the monitor feed:

- a. Specify the VLANs you want to assign to each domain.
- b. In the Unassigned VLANs column, designate a domain for unassigned VLAN traffic.

Click Help for more information.

- c. Click OK.
7. Repeat these steps to assign domains to each monitor feed.

More information:

[Managing Tenants](#) (see page 95)

Chapter 5: Managing Applications

This section contains the following topics:

[How Applications Work](#) (see page 103)

[Application Port Exclusions](#) (see page 106)

[Manage System-Defined Applications](#) (see page 113)

[Manage User-Defined Applications](#) (see page 116)

[Manage a Multi-Tiered Application](#) (see page 130)

[Application Keep-Alive Messages](#) (see page 136)

How Applications Work

An *application* specifies a TCP port or port range that you want the management console to monitor across a range of server IP addresses. For example, the management console can monitor TCP-80 traffic across a /29 server subnet.

By default, the management console automatically creates system-defined applications to monitor TCP sessions on all application ports and servers, across all client networks. For example, if the management console is configured to monitor several /24 server subnets, the management console creates a system-defined Microsoft SQL Server application to monitor TCP-1433 traffic across all servers. As new servers with matching IP addresses are provisioned, the management console automatically monitors them.

Create a port exclusion to ignore TCP sessions on a particular port or range of ports across all server subnets, a particular server subnet, or a particular server.

From the list of system-defined applications, you can create a user-defined application based on your expert knowledge rather than just the information that is available from the TCP header, to report on the actual servers that host the application. For example, to report on the performance of a particular database application, create a user-defined SQL Server application and assign an application subnet that identifies the appropriate servers which host the application.

Alternatively, if you know that the same application always runs on the same port, let the management console automatically monitor the application across all of your servers.

When registered as a data source with the CA PC or the CA NPC, the CA PC or the CA NPC automatically groups all applications for reporting.

If you have defined domains in the CA PC or the CA NPC, you can choose a domain to filter the observed application traffic.

More information:

[Managing Tenants](#) (see page 95)

[Application Port Exclusions](#) (see page 106)

How Priority Applications Work

To manage the growth of the database, if necessary, the CA Application Delivery Analysis Manager grooms the 5-minute data for low-volume applications. When the CA Application Delivery Analysis Manager grooms data, the data for the application is only available at 5-minute intervals, and the response time data in the Operations page can have a "checker board" appearance.

If you do not want the CA Application Delivery Analysis Manager to groom the 5-minute data for a system- or user-defined application, edit the Application Properties and select the Priority check box. The CA Application Delivery Analysis Manager does not groom data for a Priority application.

Find an Application

Filter the Application List to find the applications you want:

Show Subnets | Servers

Specifies an option to view observed and assigned server information:

- *Subnets* displays the corresponding server subnets where the application traffic is hosted.
- *Servers* displays the actual servers that host the application.

Search

Finds matching entries in the Application List. To clear the search results, click Clear Search.

Reset List

Removes all of your current selections in the Application List, including pages of the list that are not currently visible. If you are not sure whether an object is selected, click the Reset List command before you select the objects you want.

Show Application List

Filters system- and user-defined applications.

Max Per Page

Specifies the maximum number of application entries per page. If the list of applications fills up more than one page, the list preserves any selections you make while you navigate between pages.

Domain

(Optional) Specifies a domain to filter the observed application server port traffic.

More information:

[Managing Tenants](#) (see page 95)

[Manage User-Defined Applications](#) (see page 116)

[Manage System-Defined Applications](#) (see page 113)

Naming Conventions

The following table lists suggested naming conventions for applications.

Application Type	Suggested Naming Convention	Example
One application, single port	User-Defined- <i>portnum</i>	User-Defined-80
One application, multiple ports	User-Defined- <i>portnumStart-portnumEnd</i>	User-Defined-1024-5000 User-Defined-135-145 User-Defined-110-120 User-Defined-25-50

Application Port Exclusions

Create a port exclusion to ignore TCP sessions on a particular port or range of ports across all server subnets, a particular server subnet, or a particular server. By default, the management console creates system-defined applications to monitor all application ports on all servers that match a specified server subnet.

More information:

[Manage Server Subnets](#) (see page 72)

How Port Exclusions Work

A *port exclusion* filters the application port traffic at the monitoring device and maximizes the available resources on the management console, while at the same time focusing the management console user's attention on the applications of interest. The monitoring device ignores TCP sessions that match a port exclusion. For example, every time a user connects to a remote share, such as \\myserver\sharename, the SMB (Server Message Block) protocol opens two TCP sessions, TCP-139 and TCP-445. If the remote session is established on 445 (any Windows-based system from Windows 2000 forward), the SMB protocol will reset (RST) the 139 session and use the session established on TCP port 445. TCP-139 is used for SMB to Windows machines prior to Windows 2000. To avoid monitoring the short-lived TCP-139 sessions on all specified server subnets, create a port exclusion for port 139 and, if necessary, assign it to a domain.

Port exclusions take precedence over system- and user-defined applications. For example, if you want to create a user-defined application, and there is an existing port exclusion that matches the port range you want, edit the port exclusion to allow the management console to monitor the port range, and then create the application.

You can also use port exclusions to ignore uninteresting application server traffic that would otherwise be automatically monitored by the management console. For example, let's assume that all of your Microsoft SharePoint servers are hosted on the 192.168.43.0/24 server subnet, but 192.168.43.14 and 192.168.43.15 are test servers and you do not want to monitor them. To enable the management console to automatically monitor all of your production SharePoint servers:

- Create an application named SharePoint-80 and assign it the 192.168.43.0/24 server subnet.
- To ignore the test SharePoint servers, create a port exclusion on TCP-80 for 192.168.43.14 and 192.168.43.15.





More information:

[Managing Tenants](#) (see page 95)

How the Port Exclusion List Works

Use the Port Exclusion List to view and manage the port exclusions that specify which application ports to ignore. The Server Count and Subnet Count columns indicate the number of servers or subnets to which the port exclusion applies.

CA Application Delivery Analysis does not monitor application traffic that matches a port exclusion. You cannot create an application with a port range that matches a port exclusion.

Port Exclusion List					
<i>View and manage application ports and their assignments excluded from SuperAgent monitoring.</i>					
<input type="button" value="Add Exclusion"/>					
TCP Ports	Domain	Server Count	Subnet Count		
23	No	0	1	<input type="checkbox"/>	 
389	No	0	1	<input type="checkbox"/>	 

Add a Port Exclusion

Add a port exclusion to specify a range of ports that you want the management console to ignore across:

- All the servers in a domain. If you have not defined domains, you can configure the management console to ignore port traffic across all servers by adding a port exclusion to the default domain.
- All the servers in a server subnet. If necessary, you can create a custom server subnet to ignore application port traffic on a particular range of server IP addresses.
- One or more servers. The management console can ignore port traffic for a particular server or servers. If necessary, you can add a server to the port exclusion.

When adding a port exclusion, if there are existing applications that match the port exclusion, the management console stops monitoring those applications.

When deleting an application, you can optionally create a port exclusion to prevent the management console from subsequently automatically monitoring the application.

Follow these steps to add a port exclusion:

1. Click the Administration page.
2. Click Data Monitoring, Applications in the Show Me menu.
3. (Optional) Choose the domain to which you want to add the port exclusion. If you have not defined domains to separate duplicate IP traffic, this option excludes port traffic across all monitoring devices in the default domain.
4. Scroll to the Port Exclusion List at the bottom of the page and click Add Exclusion. Port Exclusion Properties opens.
5. Specify the Starting Port and Ending Port to exclude a particular port or a range of ports.
The Starting Port must be the same or lower than the Ending Port.
6. (Optional) To exclude the port range across all servers in the currently selected domain, choose Ignore application port traffic on all servers in the domain. Note that if you have not defined domains, the default domain applies to all servers.
When you select this option, the management console prompts you to confirm that you want to ignore port traffic on the domain.
7. (Optional) To exclude the port range across a range of servers:
 - a. Click Assign Subnets.
 - b. In Port Exclusion Subnets, assign an existing server subnet or create an application subnet to exclude port traffic across a range of servers.
To create an application subnet, specify the IP address and mask, then click Add Application Subnet.

To assign an existing server subnet, double-click the subnet you want. The list of available subnets includes the server subnets that the management console uses to monitor the servers in your environment, and any application subnets you have created.

- c. Click Apply.
8. (Optional) To exclude the port range across a particular server or servers:
 - a. Click Assign Servers.
 - b. In Port Exclusion Servers, double-click an available server to add it to the port exclusion.
 - c. Click Apply.
9. Click OK.
10. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

More information:

[Managing Servers](#) (see page 67)


[Managing Tenants](#) (see page 95)

[Delete a User-Defined Application](#) (see page 129)

Edit a Port Exclusion

Edit a port exclusion to update the range of ports and servers that you want the management console to ignore.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Applications in the Show Me menu.
3. (Optional) Choose the domain where you want to edit a port exclusion.
4. Scroll to the Port Exclusion List at the bottom of the page, and click  to edit a port exclusion.

Port Exclusion Properties opens.

5. Specify the port exclusion properties and click OK.

For information about specifying port exclusion properties, click Help.

6. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.


More information:


[Managing Tenants](#) (see page 95)

Delete a Port Exclusion

Delete a port exclusion to resume monitoring of the previously excluded port traffic.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Applications in the Show Me menu.
3. (Optional) Choose the domain from which you want to delete the port exclusion.
4. Scroll to the Port Exclusion List at the bottom of the page, and click  to delete an exclusion.

(Optional) To delete more than one application exclusion, select the appropriate exclusions, and click  to delete all selected exclusions.

5. Click Continue with Delete at the prompt to let the management console begin to automatically monitor matching application port traffic.

The exclusion is deleted.

6. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

More information:

[Managing Tenants](#) (see page 95)

Manage System-Defined Applications

The management console creates system-defined applications, based on the list of client networks and server subnets you specify, to monitor the busiest Standard applications and active FTP (TCP-20 and TCP-21) applications on each matching server.

When possible, let the management console automatically monitor application traffic across all of your servers. Unlike a user-defined application, servers cannot be assigned to a system-defined application. If the management console requires additional information beyond what is available in the TCP packet to monitor the application, create a user-defined application. For example, create a user-defined application to monitor:

- An application that communicates on a range of ports
- A particular Web Application that communicates on TCP-80

You cannot set a performance OLA or an availability OLA for a system-defined application. However, if you want to apply an OLA to all the servers monitored by the management console, create a user-defined application and assign it all servers in a domain.

To reduce the number of application ports that are automatically monitored by the management console:

- Delete an application.
- Create a port exclusion.

More information:

[Delete a System-Defined Application](#) (see page 115)

[Application Port Exclusions](#) (see page 106)

[Assign Servers to an Application](#) (see page 126)

Edit a System-Defined Application

Edit a system-defined application, for example, to:

- Rename the application. By default, the management console names System applications on well-known ports.
- Prioritize the application and prevent the management console from grooming or filtering the application.
- Change the default incident response. Note that the default application incident response does not specify any responsive actions.

The status of a System-defined application does not change.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Applications in the Show Me menu.
3. If you have defined domains in the CA PC or the CA NPC, you do not need to choose a domain. Any changes you make to system-defined application properties apply across domains.
4. Browse the Configured By column in the Application List to find a System-configured application.

If necessary, click Reset List to reorganize the Application List.

To edit more than one System-defined application, select the applications you want. Note that you cannot rename more than one application with the same name.

5. Select the application you want and click Edit.

The Application Properties opens.

6. Complete the fields in Application Properties and click OK.

For information about setting application properties, click Help.

7. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

More information:

[Managing Tenants](#) (see page 95)

[Manage User-Defined Applications](#) (see page 116)

Delete a System-Defined Application

Delete system-defined applications to optimize the available system resources on the management console and its monitoring devices. You can prevent the management console from automatically monitoring an application after you delete it by creating a port exclusion.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Applications in the Show Me menu.
3. If you have defined domains in the CA PC or the CA NPC, you do not need to choose a domain. Any changes you make to the application apply across domains.
4. Browse the Configured By column in the Application List to select a System-configured application and click Delete.

If necessary, click Reset List to reorganize the Application List.

To delete more than one System-defined application, select the applications you want. Note that you cannot rename more than one application with the same name.

5. Choose an option to delete the application:

Delete and Add Port Exclusions

Delete the selected applications and prevent the management console from automatically monitoring the application.

Delete

Delete the selected applications, but enable the management console to automatically monitor the application if it sees matching application traffic.

6. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

More information:

[Managing Tenants](#) (see page 95)

[Application Port Exclusions](#) (see page 106)

Manage User-Defined Applications

When possible, let the management console automatically monitor application traffic across all of your servers. If the management console requires additional information beyond what is available in the TCP packet to monitor the application, create a user-defined application and assign particular servers or a range of server IP addresses. For example, create a user-defined application to monitor:

- An application that communicates on a range of ports
- TCP-80 traffic on a particular server or servers

Use the list of system-defined applications to identify your busiest application ports and then leverage your application expertise to create user-defined applications to monitor your most heavily-trafficked, business-critical, and time-sensitive TCP applications.

When creating a user-defined application to monitor application server traffic that is currently monitored by a System application, keep in mind that after you create a user-defined application:

- The System application stops reporting on matching servers that are assigned to the user-defined application. If the management console observes application traffic on a server that is not identified by the user application's server assignment, the management console reports the application server response time in the System application.
- The management console collects new data to report on the user-defined application. The management console does not report the existing System application data in the User application.
- Port exclusions take precedence over system- and user-defined applications. If necessary, remove the port you want from any port exclusions and then create a user-defined application.

More information:

[Application Port Exclusions](#) (see page 106)

[Assign Servers to an Application](#) (see page 126)

Create a Standard Application

A *Standard application* is a typical TCP application that connects to a port on the server. All traffic flows between that port and a port on the client. Any type of monitoring device can monitor a Standard application.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Applications in the Show Me menu.
3. (Optional) Choose a domain to filter the list of available server subnets and servers.
4. Create a user-defined application using one of the following methods:

Create a new application

Click Create New Application.

Create a new application from an existing system-defined application

Click a System application and then click Create New Application. To create an application that communicates on a range of ports, click the System applications you want.

5. Set the Application Type to Standard and then complete the fields in Application Properties:
 - Application Name.
 - Make this a Priority Application. Choose this option to specify that you do not want the management console to groom or filter the application.
 - Beginning Port. Beginning TCP port number for the port range.
 - Ending Port. Ending TCP port number for the port range.
 - Port Side. Choose an option to specify how the Standard application responds to client requests:
 - Application listens on these ports. Choose this option when the server listens for client requests within the specified port range. This is the default.
 - Application talks to these ports. Choose this option when the server responds to client requests within the specified port range on the client. This option is not applicable when monitoring an application with a Cisco NAM monitoring device.
 - Incident Response. Choose an application incident response to specify how Application Delivery Analysis responds to a Network or Server incident that affects the application.

- **Availability Monitoring.** Choose an option to enable or disable availability monitoring on the application. When enabled, the management console passively observes the availability of the application, and, if necessary, actively checks its availability.
- **Notes.** (Optional) additional information about the application.

For information about setting Standard application properties, click Help.

6. Click Next to assign server subnets and servers to the application, then click OK.

For information about assigning servers to an application, click Help.

7. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

More information:

[Managing Tenants](#) (see page 95)

[Assign Servers to an Application](#) (see page 126)

Create a Web Application


Create a Web application to report on the TCP response time of HTTP traffic. To monitor HTTP traffic, you must use a CA Standard Monitor.

Specify the resource path of the Web application that you want to monitor. The resource path must match what is in the HTTP Request header (GET, POST, HEAD, or TRACE). The management console uses the resource path to identify particular HTTP traffic on a Web server.

Important! If the management console only reports on the *appname-Other* application, verify by packet capture that the resource paths which you have specified match the HTTP Request header. In some cases, such as when monitoring web traffic that is traversing a proxy, it may be necessary for you to specify the full URL, for example, *http://server/resource*, rather than the resource path as shown in the examples below:


/SuperAgent

Identifies all HTTP traffic to CA Application Delivery Analysis. For example:



/pc

Identifies all HTTP traffic to the CA NPC. For example:



The management console creates a separate Web application for each resource path you define, and a *appname-Other* application to monitor all other HTTP traffic across all the servers that are assigned to the application. Use the *appname-Other* application to analyze resource changes on the Web site. For example, if the *appname-Other* application experiences a performance degradation, analyze the HTTP traffic on the server and, if necessary, add a resource path.

If you have a load-balanced Web application, assigning a server subnet to the application enables the management console to automatically monitor the response time of each resource path as servers are provisioned.

More information:

[Internet-Facing Web Applications](#) (see page 121)

Web Application Considerations

When creating a Web application, keep in mind the following:

- You must use the Packets feed, a CA Standard Monitor, to monitor a Web application. The CA Standard Monitor is the only monitoring device that monitors a Web application. After you create a Web application, verify the Packets monitor feed is assigned to each server that belongs to the Web application.
- Do not create a Web application to monitor all HTTP traffic on a server. If you want to monitor all TCP-80 or 8080 traffic on a server, create a Standard application. A Web application is a monitoring device-intensive process because the monitor must process the HTTP header.
- Do not monitor Web applications on more than two non-standard ports. A non-standard port is a port other than TCP-80 or TCP-8080. For example, you can monitor Web applications on TCP-80 and TCP-8080, and two other ports. The additional resources that are required to process non-standard ports can negatively impact management console performance.
- When a user accesses a Web application through a proxy server, the management console reports the application traffic from the proxy server's client network.
If the proxy server uses X-Forwarded-For (XFF), the management console can translate the XFF header to report on the actual client network rather than the proxy server's client network.
- The management console does not report HTTP sessions from any Web applications.
- The management console cannot monitor Web application traffic on HTTPS (TCP-443) because the URLs are encrypted. All such traffic shows up in the "Other" application.

More information:

[Edit a Server](#) (see page 79)

[Support for XFF Translation](#) (see page 270)

Internet-Facing Web Applications

When making response time calculations, the management console assumes that it is positioned next to the application server. Monitoring response time close to the server enables the management console to accurately measure network and server response time.

Monitoring a web site on the internet will result in skewed metrics because the management console is next to the client, not the server. Server Response Time measured at the client will include the network latency to the server.

When monitoring an internet-facing web application, we recommend that you disable performance thresholds on network performance metrics and only monitor performance threshold status for server performance metrics. Thresholding for network metrics makes sense for client subnets which are specific to a particular site and which share a particular circuit. The internet is too broad a category for meaningful network performance thresholds.

If you want to monitor a third-party service like google.com, use an IPSLA test to record response time; however, there is no way to know when a spike is due to Internet latency, server response, or an application issue.

More information:

[Edit Performance Thresholds](#) (see page 150)

Create a Web Application

Create a Web application by specifying the path of the resource on the Web server that you want to monitor. The management console automatically creates a Web application to report on matching HTTP traffic across all servers that are assigned to the application.

Important! To monitor HTTP traffic for a specific resource, you must use the Packets monitor feed on a CA Standard Monitor.

You must specify at least one resource path that you want to monitor.

Important! If you want to monitor *all* HTTP traffic on a server, create a Standard Application that monitors TCP-80 or TCP-8080 traffic.

Follow these steps to create a Web Application:

1. Click the Administration page.
2. Click Data Monitoring, Applications in the Show Me menu.
3. (Optional) Choose a domain to filter the list of available server subnets and servers.
4. Scroll to the Application List and click Create New Application.
5. Set the Application Type to Web.
6. Complete the fields in Application Properties.

For information about setting Web application properties, click Help.

7. Click Next to assign server subnets and servers to the application, then click OK.

For information about assigning servers to an application, click Help.

8. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

9. Edit the server properties to verify that each server which is assigned to the application is also assigned to the Packets monitor feed on a CA Standard Monitor. The Packets monitor feed is required to monitor a Web application.

More information:

[Managing Tenants](#) (see page 95)

[Edit a Server](#) (see page 79)

[Assign Servers to an Application](#) (see page 126)

Create an FTP Application

Create an FTP application to monitor active FTP on a particular server. The management console automatically monitors all active FTP sessions. All types of monitoring devices can monitor an FTP application.

In active FTP, the management console monitors the command and data ports as a pair to determine response time because the request and response actions take place on different ports.

In active FTP:

TCP Port	Is the
20	Data port. The FTP server's local data port to transfer data to the client.
21	Command port. A client connects to this port to send FTP commands.

To monitor an active FTP application that uses different ports, create a Control Port application.

Note that the management console also automatically monitors passive FTP. However, because passive FTP opens a random, unprivileged port on the server to transfer data to the client, the application activity is low-volume. Therefore, it is unlikely that a passive FTP application will bubble-up into the Operations page of the management console.

To create an active FTP application follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Applications in the Show Me menu.
3. (Optional) Choose a domain to filter the list of available server subnets and servers.
4. Scroll to the Application List and click Create New Application.
5. Set the Application Type to FTP.
6. Complete the fields in Application Properties.

For information about setting FTP application properties, click Help.

7. Click Next to assign server subnets and servers to the application, then click OK.

For information about assigning servers to an application, click Help.

8. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

More information:

[Managing Tenants](#) (see page 95)

[Create a Control Port Application](#) (see page 125)

[Assign Servers to an Application](#) (see page 126)

Create a Control Port Application

Create a Control Port application to monitor applications where the Control port sends and receives the request information, and the Data port sends and receives the actual data. The management console must monitor both ports to determine the transaction response time. All types of monitoring devices can monitor a Control Port application.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Applications in the Show Me menu.
3. (Optional) Choose a domain to filter the list of available server subnets and servers.
4. Scroll to the Application List and click Create New Application.
5. Set the Application Type to Control Port.
6. Complete the fields in Application Properties.

For information about setting Control Port application properties, click Help.

7. Click Next to assign server subnets and servers to the application, then click OK.

For information about assigning servers to an application, click Help.

8. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

More information:

[Managing Tenants](#) (see page 95)

[Assign Servers to an Application](#) (see page 126)

Assign Servers to an Application

Use the Application List to view and manage the servers that host an application.

If the servers that host an application have a contiguous range of IP addresses, assign a server subnet rather than a particular server to an application. Specifying a server subnet optimizes management console performance and enables automatic server assignment.

Reuse your existing server subnets or create an application subnet that more closely aligns with the actual servers that host an application.

An *application subnet* collects data across a range of server IP addresses, and its subnet mask can be wider or narrower than your existing server subnets. You can reuse an application subnet by assigning it to more than one application.

We recommend assigning a server subnet or application subnet to an application. This approach works best when your server administrator provisions application servers using contiguous IP ranges. Avoid assigning servers to an application. If the host resources change, you must update the application's server assignment to continue monitor the application.

Assigning a:

- Domain

Enables the management console to automatically assign all matching servers in the domain to the application, and keep the application server assignments up-to-date as your server subnets change.

If you have not created domains to separate traffic, assigning the default domain assigns all servers monitored by the management console.

- Range of server IP addresses specified by an existing server subnet or application subnet.

For example, if you configured the management console to monitor servers using server subnets that closely align with your server VLAN definitions, you can assign the same server subnet to an application.

- New application subnet.

For example, if you have an existing /22 server subnet, but for a particular application, you know its /26 server subnet definition, create the /26 application subnet and assign it to the application.

- Server.

You can assign a server from the Server List, or add a new server.

Follow these steps:

1. Click the Administration page.

2. Click Data Monitoring, Applications in the Show Me menu.
3. (Optional) Choose a domain to filter the list of available server subnets and servers.
4. Scroll to the Application List, select an application and click Edit.
The Application Properties opens.
5. Click Assignments to assign servers to the application.
For information about assigning servers to an application, click Help.
6. (Optional) Deselect a server subnet, application subnet, or server to unassign it.
Note that the console automatically deletes an application subnet when it is not assigned to any applications.
7. If application availability monitoring is enabled, and a load-balancer distributes the application traffic between servers, to enable the management console to check server availability, specify the number of servers that must be available.
8. Click OK.
If the OK button is disabled, make sure you have assigned a server or subnet to the application, and click Properties to verify the application properties are properly specified.
9. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.
Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

More information:

[Managing Tenants](#) (see page 95)

[Check Server Availability](#) (see page 206)

Edit a User-Defined Application

Edit the properties of a user-defined application, for example, to add a URL to a Web application.

If you edit more than one application at the same time, you can specify common application properties, such as availability monitoring.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Applications in the Show Me menu.
3. (Optional) Choose a domain to filter the list of available server subnets and servers.
4. Scroll to the Application List and select the applications you want to edit, then click Edit.

If necessary, click Reset List to remove any selections from the Application List.

5. Click Properties to edit the application settings.

For information about application properties, click Help.

6. Click Assignments to edit the server assignments for the application and click OK.

For information about assigning servers to an application, click Help.

7. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

If you are not prompted, the changes to the application properties did not require you to synchronize monitoring devices.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

More information:

[Managing Tenants](#) (see page 95)

Delete a User-Defined Application

Delete a user-defined application to remove it from the Application List, and optionally, create a port exclusion to prevent the management console from attempting to automatically monitor the corresponding application port.

When considering which applications to delete, keep the following points in mind:

- Avoid monitoring applications that are not time-sensitive, such as a backup application where the response time of the application is less critical.
- Applications that take more system resources on the monitoring device and the management console to process rather than filter out, and are less critical, are good candidates to delete. For example, a backup application can generate traffic across all the client IPs in your environment, which can consume a lot of rows in the database, and create a higher load on each monitoring device.

After you delete an application, existing data continues to be available for reporting purposes according to your database settings.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Applications in the Show Me menu.
3. If you have defined domains in the CA PC or the CA NPC, you do not need to choose a domain. Any changes you make to the application apply across domains.
4. Scroll to the Application List, select an application from the list and click Delete.

The management console prompts you to confirm the deletion and optionally, prevent the management console from automatically monitoring the application by adding a matching port exclusion.

- Click Delete to delete the application and allow the management console to recreate the application from observed application traffic.
 - Click Delete and Add Port Exclusions to delete the application and create a port exclusion rule to prevent the management console recreating the application from observed application traffic.
5. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

More information:

[Managing Tenants](#) (see page 95)

[Edit Database Storage Preferences](#) (see page 221)

Manage a Multi-Tiered Application

Use the management console to gain visibility into the network, server, and application performance of each tier of a multi-tiered application. A *multi-tier application* is an application with more than one server, and communication between servers is performed by at least one server that acts as both a server to client requests, and a client to another server.

More information:

[How Multi-Tiered Applications Work](#) (see page 131)

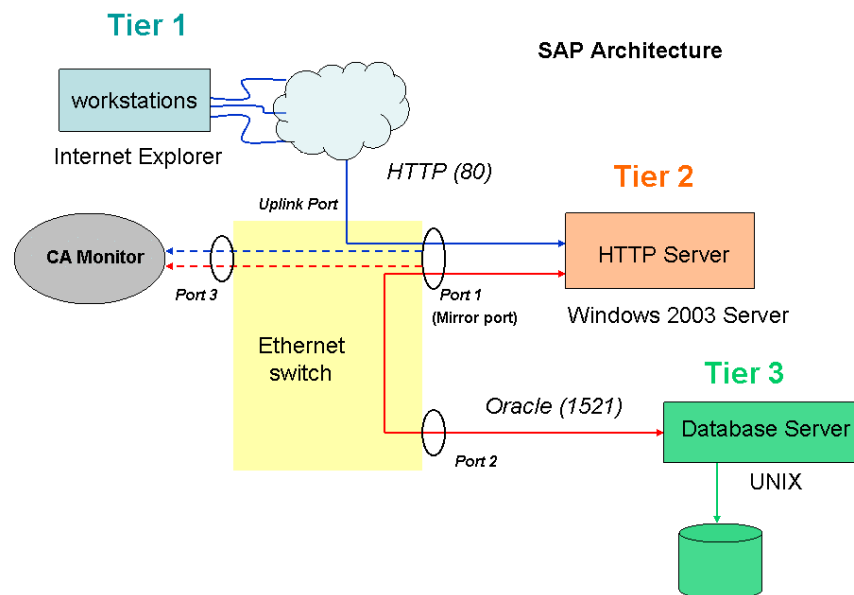
[How To Monitor a Multi-Tiered Application](#) (see page 132)

How Multi-Tiered Applications Work

Consider an N-Tier SAP architecture that consists of the following tiers:

- Tier 1--Internet Explorer running on a user workstation
- Tier 2--An HTTP-based application running on Windows
- Tier 3--A database server running Oracle on UNIX

In a multi-tier application, at least one server acts as both a server and a client to other application servers. In the example above, Tier 2 is a server to user requests from Tier 1, and a client to requests from the Tier 3 server.



The following process occurs:

1. Using Internet Explorer, a user initiates a connection to the Tier 2 HTTP server, which is illustrated by the blue line.
2. After the connection is established, the user requests application data.
3. The HTTP server forwards this request to the Tier 3 Oracle database server, depicted by the red line.
4. The Oracle server runs the user query and returns the results to the Tier 2 HTTP server.
5. The HTTP server sends the data back to the Tier 1 client.

The multiple handoffs among the application tiers can make it difficult to identify the source when a performance problem occurs with an N-Tier application. Operationally, when Tier 2 waits for the Tier 3 response, its performance depends on the Tier 3 performance.

How To Monitor a Multi-Tiered Application

Complete these steps to configure the management console to monitor and report on an N-Tier application architecture:

- Mirror the application conversations of interest to a monitoring device.
- Configure the management console to monitor the application architecture.

More information:

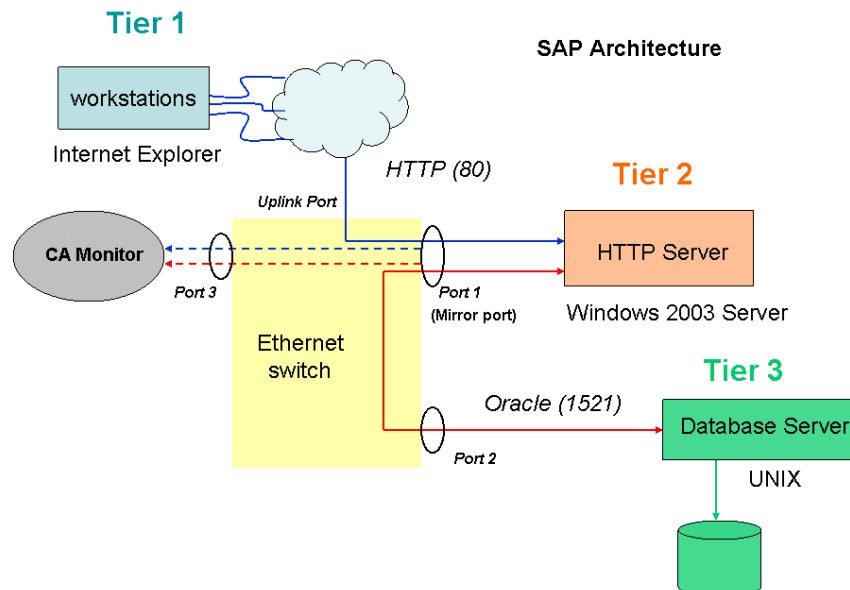
[Strategy for Mirroring Application Conversations](#) (see page 133)

Strategy for Mirroring Application Conversations

For the management console to gain visibility into an N-Tier application, you must mirror host conversations between each tier to a monitoring device. To start this process, run mirroring commands to the Ethernet switch on which the servers and the monitoring device are connected. The mirror commands cause the packets to and from the various hosts to be copied to the monitoring device. This copy function has minimal performance impact on the switch.

If the multi-tier application is virtualized, consider mirroring the server-to-server traffic to the virtual switch on the ESX host, and mirroring the front-end server traffic to a physical monitor.

The following figure depicts switch port mirroring.



In complex environments with more than two servers, you might need to mirror several ports to capture all tiers of an application architecture. Carefully select the ports to mirror in this environment to avoid capturing the same conversation twice.

More information:

[Monitoring Device Recommendations](#) (see page 251)

Create a Multi-Tier Application

To monitor a multi-tier application, create an application for each tier, and use a naming convention to help you more easily identify and report on some or all the tiers.

Use a naming convention that facilitates administration, reporting, and analysis so that a management console user can immediately recognize that a dependent relationship exists between each application tier. Typically, the performance of an application tagged as Tier 2 often depends on the performance of an application tagged as Tier 3. Review the Tier 3 application performance when you analyze Tier 2 application performance.

The following table shows an example of a multi-tiered application. If you define each tier of the applications this way, each application appears next to each other in the management console. This method shows the multiple pieces of the application architecture and reminds you that there is a dependent relationship among various elements of applications and processes:

Application name	Start port	End port	Port side	Associated server
SAP-HTTP-(80)-Tier 2	80	80	Application listens on these ports	HTTP
SAP-Oracle-(1521)-Tier 3	1521	1521	Application listens on these ports	Oracle

Follow these steps:

1. Add the Tier 1 client networks to the Network List and be sure to specify a 24-bit (or higher) subnet mask.
2. Add all servers that participate in the application tier to the Server List and verify that the correct monitor feed is associated with the server.

When you add a server, it is automatically added to the Network List as a host with a 32-bit mask, which indicates that the server acts as a client as is the case in N-Tier architectures. In the previous example, the HTTP server acts as a client to the Oracle database server.

3. Add the applications to the management console. Use the following naming convention for N-Tier applications:

`<ApplicationName>-<Protocol/Function>-(<TCPPort>)-<Tier#>`

where the variables are defined as follows:

<ApplicationName>

Is the name of the application.

<Protocol/Function>

Is the application daemon running on the server.

<TCPPort>

Is the daemon port number.

<Tier#>

The Tier number.

4. Repeat these steps to define each application tier.
5. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

More information:

[Managing Client Networks](#) (see page 29)

[Managing Servers](#) (see page 67)

[Manage User-Defined Applications](#) (see page 116)

Application Keep-Alive Messages

Periodic keep-alive messages exchanged between clients and servers are common in IP networks and let servers determine whether clients are still active or reachable. The management console generally ignores TCP keep-alives if they adhere to the standards in RFC 1122 and excludes any associated statistics from byte counts and observation totals.

However, some applications are designed to use custom keep-alive mechanisms. If the response from the client is an acknowledgment that contains a payload, the management console treats the client response as a request for data, and starts the Server Response Time (SRT) timer. This results in an inaccurate SRT measurement and SRT observation count once the keep-alive timer expires, when the server typically sends out another packet.

Popular applications that use keep-alives include Citrix and Microsoft Exchange. If you suspect that another application is sending keep-alives, look for the inverse relationship between observations and SRT and for SRT averages in the second range instead of the millisecond range.

A CA Standard Monitor or CA Multi-Port Monitor can be configured to filter application keep-alive messages by Server Response Time to avoid skewing server metrics.

The management console uses NRTT observations to filter applications that use application keep-alives. If necessary, you can adjust the threshold for the minimum number of NRTT observations during a 5-minute interval.

More information:

[Manage Console Settings](#) (see page 223)

[Filter Out Keep-Alive Messages](#) (see page 281)

Chapter 6: Managing Performance Thresholds

This section contains the following topics:

[How Performance Thresholds Work](#) (see page 138)

[How Incidents Open and Close](#) (see page 147)

[Edit Performance Thresholds](#) (see page 150)

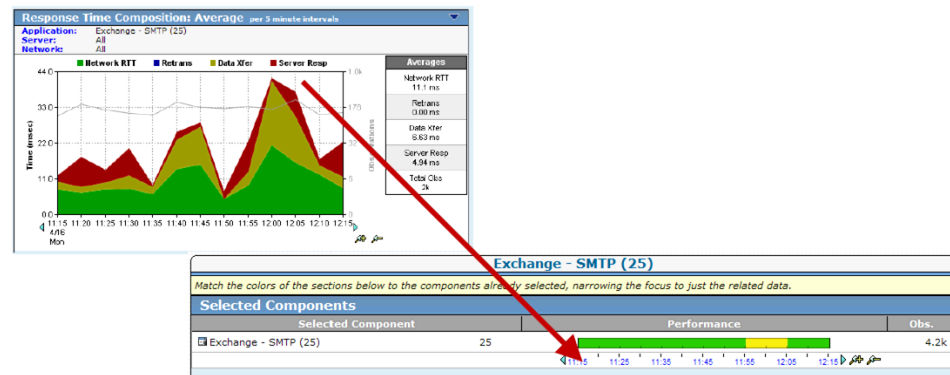
[Add Performance Thresholds](#) (see page 154)

[Enable Default Performance Thresholds for a Group of Networks](#) (see page 155)

[Edit Performance Thresholds for WAN-Optimized Network Segments](#) (see page 156)

How Performance Thresholds Work

Performance thresholds turn an advanced trend report from the Engineering page into a bar chart report on the Operations page that is easily read.



Performance thresholds help you recognize performance problems and allow the management console to automatically launch responses to notify or investigate a problem.

Thresholds are boundaries of acceptable performance behavior that exist for each system- and user-defined application. Thresholds are important because they:

- Enable the management console to rate data.
- Contribute to incident creation and the resultant incident responses and investigations, which enables timely troubleshooting and problem resolution.

When performance thresholds are properly configured, the colors on the bar chart correspond to real performance issues or metric degradations. Accurate threshold tuning makes pinpointing the cause of these issues a simple matter of matching colors.

The yellow and orange severity indicators are intended to draw attention to Minor and Major performance degradations. We recommend calibrating threshold settings to make sure the severity indicators correspond to the actual conditions that cause network users to submit Help Desk tickets.

When the management console sees degraded application performance on a server or network, the management console automatically opens an incident. Available from the Incidents page, an *incident* creates a record of information about a performance problem.

After your performance threshold settings are properly adjusted, choose how the management console should respond to network and server incidents, for example, by sending an email notification when application performance on a particular client network degrades.

More information:

[Managing Incident Responses](#) (see page 165)

[How Incidents Open and Close](#) (see page 147)

How Application Performance is Rated

The management console rates the performance of an application by observing its TCP transactions and calculating:

- Network metrics for each client network that communicates with the application. If the 5-minute average for a Network metric exceeds the threshold, and the management console observed the metric the minimum number of times, the management console rates the corresponding 5-minute interval for the client network as Minor (yellow) or Major (orange) and creates a network incident.
- Server metrics for each server that hosts the application. If the 5-minute average for a Server metric exceeds the threshold, and the management console observed the metric the minimum number of times, the management console rates the corresponding 5-minute interval for the server as Minor (yellow) or Major (orange) and creates a server incident.
- Combined metrics for the application itself, which include both network and server metrics. If the 5-minute average for a Combined metric exceeds the threshold, and the management console observed the metric the minimum number of times, the management console rates the corresponding 5-minute interval for the application as Minor (yellow) or Major (orange).

Note that the management console does not create application incidents. However, because Combined metrics include both Network and Server metrics, the management console can rate a server or network as Minor (yellow) or Major (orange), and rate the corresponding performance impact on the application itself. For example, if a Server metric degrades, the management console can also rate a Combined metric for the application as Minor.

To rate performance data as Normal, Minor (yellow), or Major (orange), the management console must collect 2 full business days of data, counting a business day from GMT midnight to midnight. For example, if the management console begins collecting data for a TCP session between a server port and client network on Monday at 3:30 p.m. EST, the management console cannot rate the performance of the application on that network until 7:00 p.m. EST on Wednesday. If the management console has not collected 2 full business days of data, the management console rates the TCP sessions between the server port and client network as Unrated.

By comparison:

- Baselines, which are available from the Explore button on the Operations page, report the historical norm for all TCP sessions between an application port on a server, and a client network. The management console calculates hourly baselines between all application ports, servers, and client networks, and tracks them for day of week, day of month, and activity of the past week. The management console uses the baselines to indicate when a performance condition is normal for that hour of the day. The management console requires 2 full business days of data to calculate baselines.

- Operational Level Agreement (OLA) reports, which are available from the Management page, quantify current performance and performance trends. Performance OLAs show how well an application is performing by counting, on an hourly basis, the percentage of transactions that are faster than a particular threshold.

More information:

[Managing Application Performance OLAs](#) (see page 189)

[Combined Metrics](#) (see page 143)

[Network Metrics](#) (see page 142)

[Server Metrics](#) (see page 143)

How Performance Metrics Work

The following sections describe the metrics that the management console uses to rate application performance. If the application is WAN-optimized by Cisco WAAS or Riverbed Steelhead, not all the performance metrics apply to the Client, WAN, and Server segments.

The management console calculates the following metrics from a standard TCP transaction:

- Network Metrics
- Server Metrics
- Combined Metrics

Network Metrics

Network metrics indicate an application performance problem is caused by a network that is communicating with the application.

Network Round Trip Time

Is the amount of time that a packet takes to travel between the server and clients on a network, excluding loss. Application, server, and client processing time are excluded.

Network Connection Time

Is the amount of time between the SYN-ACK sent by the server and the ACK received back from the client. When a network is uncongested, it is a measurement of network latency that represents the minimum latency due to distance and serialization, and is the best possible round trip time for your network architecture.

Sudden spikes in this value are commonly attributed to congestion, while a plateau (which goes up and stays up) typically indicates a path change.

Effective Network Round Trip Time

Consists of Network Round Trip Time plus Retransmission Delay. Note that Retransmission Delay is not the delay due to any retransmissions; it is the average amount of retransmission delay per round trip. It is important to note that the management console is adding two averages, and is actually combining two metrics.

Retransmission Delay

Is the elapsed time between the original packet send and the last duplicate packet send. The management console reports Retransmission Delay as an average across observations and not just for the retransmitted packets. If one packet in a set of 10 requires 300 ms of retransmission time, the Retransmission Delay is reported as 30 ms (300 ms/10 packets).

Server Metrics

Indicate an application performance problem is caused by a server that hosts the application.

Server Response Time

Is the time it takes for a server to send an initial response to a client request or the initial server "think time." Increases in the Server Response Time generally indicate the following:

- A lack of server resources such as CPU, memory, disk, or I/O
- A poorly written application
- A poorly-performing tier in a multi-tier application

Server Connection Time

Is the amount of time that a server takes to acknowledge the initial client connection request by sending a Syn-Ack in response to the client's SYN packet.

Refused Session Percentage

Is the percentage of connection requests that were explicitly rejected by the server during the three-way handshake (see definition on page 423). Part of the Unfulfilled TCP/IP Session Requests report.

Unresponsive Session Percentage

Is the percentage of sessions where a connection request was sent, but the server never responded. Part of Unfulfilled TCP/IP Session Requests report.

Combined Metrics

Indicate an application performance problem is caused by both a server that hosts the application and a network that is communicating with the application.

Transaction Time

Is the time it takes to transmit a complete application response measured from the first response (the end of the Server Response Time) to the last packet sent in that request and that it can be impacted by the design of the application, the performance of the server or the network.

The management console shows this type of response time data in the Response Time Composition: Average report on the Engineering page. The management console does not open incidents when the Transaction Time threshold is crossed.

Data Transfer Time

Is the elapsed time between when the server starts responding and it finishes sending data. Factors such as the response sizes, the bandwidth available on the network, and interaction between the application and the network affect the value.

The management console does not open an incident when the Data Transfer Time threshold is crossed.

Options to Customize Performance Thresholds

Specify a performance threshold value by:

- Entering a Sensitivity level. The management console dynamically generates a threshold value based on the Sensitivity level.
- Entering a static value in milliseconds or specify a percentage
- Disabling a metric threshold

By default, the Sensitivity option is enabled. No matter which method is used, the threshold always includes a value for the minimum number of observations. An *observation* is an opportunity for the management console to calculate a metric from a TCP transaction. To rate the performance of a metric as Normal, Minor (yellow), or Major (orange), a monitoring device must see a minimum number of observations.

Alternatively, you can choose to not include a particular metric when rating the performance of a network, server, or application.

Sensitivity Levels (Dynamic Values)

When you specify a Sensitivity level, the management console automatically generates a new threshold value for a metric each night at GMT midnight, using percentile statistics from the last 30 days to determine an appropriate setting. The management console generates a separate set of threshold values for the users who access an application from each client network.

As a result, the management console does not apply the same threshold value to users on high-latency remote links as it does to users on a local LAN, and a threshold value represents historically extreme performance. For example, consider an application server located at a data center in Miami, with users who access the application locally and from a remote site in Munich, Germany.

Performance by Network			
Network	Subnet	Network Round Trip Time	Observations
		Wtd. Average: 12.28 ms	Average: 32.59 ms
Munich	192.168.25.0/24	140.67 ms	406,988
Miami	192.168.6.0/24	5.68 ms	1,251,756

If you use Sensitivity levels to set performance thresholds for this application, the management console generates a separate set of threshold values for the users who access the application from each client network. Specifying Sensitivity levels also means that you are not required to define the same application twice, such as Exchange to Munich and Exchange to Miami, to set different thresholds for application performance across remote locations.

When the sensitivity of a threshold is raised, the threshold value is lowered because the intent is to receive more incidents based on performance levels. The management console sensitivity threshold is similar to pain sensitivity. People who are highly sensitive have a low threshold for pain, and will cry out more often. People who are less sensitive can handle more pain and are much less likely to cry out.

Sensitivity can be adjusted from a value of 0 (insensitive) to 200 (very sensitive):

- A sensitivity setting of 200 results in rating traffic measuring over the 75th percentile as Minor (yellow) or Major (orange).
- A lower sensitivity setting yields a higher threshold and results in fewer incidents.

The Sensitivity level does not change when the management console recalculates threshold values. To view today's threshold value for a Sensitivity level, use the Sensitivity calculator.

More information:

[Edit Thresholds from the Administration Page](#) (see page 151)

Static Thresholds (Static Values)

Selecting the Milliseconds or Percentage threshold setting lets you specify a static threshold value for a particular metric. This value does not change unless you change it.

Static thresholds can be used to define target values for metrics that are highly consistent or that have known values for poor performance. For example, a healthy server should not refuse session requests from clients, so setting a static threshold of 1% refusals for Minor (yellow) and 3% for Major (orange) would prevent the management console from ignoring server issues on a server that frequently refuses sessions.

Other metrics that are good candidates for static thresholds are:

- Unresponsive sessions. Should be zero on healthy servers. Set low thresholds of 5 to 10%.
- Server Connection Time. Should be sub-millisecond. Set to alarm at 2 to 5 ms as long as the management console is collecting data on the same switch as the application server.

Using the Milliseconds or Percentage static-value setting also means that thresholds are no longer automatically separated out by network and by server. To resolve this issue, you must group networks with similar latency by network type and then assign a custom set of thresholds to each network type.

Before configuring a static threshold for a given application, view a report on the Engineering page, and use the Settings button to select the metric in question. Make sure you are not going to set any servers, networks, or applications at a permanently degraded threshold level without first taking note of their metrics and correcting any underlying issues.

More information:

[Add Performance Thresholds](#) (see page 154)


How Incidents Open and Close

The management console opens network and server incidents when the 5-minute average for a network or server metric exceeds the threshold.

The management console automatically closes an incident when:

- One full clock hour (from the top of hour to the top of hour) of acceptable performance behavior has transpired.
- The server that is associated with the network or server incident enters a maintenance window. Note that after a scheduled maintenance period ends, the management console can open a new incident, if the performance problem persists.
- The incident is 24 hours old. Note that the management console opens a new incident if the performance problem persists.

You can also Acknowledge an incident to indicate your awareness of the issue. In the following example, the management console opened a network incident on the 191.168.1.0/24 network at 7:05 a.m. After 7:20 a.m., there were no subsequent network threshold violations, so the management console closed the incident at 9:00 a.m. Note that because the management console closes incidents at the top of hour, when you Acknowledge an incident, the report also marks the hour (from the top of hour) before the incident, which is 7:00 a.m., to show the time frame of the incident:



If you have set OLAs for availability, the management console opens a server incident when an application is rated Unavailable.

More information:

[Managing Application Availability](#) (see page 201)

[Schedule Server Maintenance](#) (see page 87)

[How Application Performance is Rated](#) (see page 140)

NetQoS Performance Center (CA NPC)

When the management console is registered as a data source with the CA NPC, and the management console sees a new Minor (yellow) or Major (orange) condition, the management console opens an incident and CA NPC opens a corresponding event.

- If the Minor or Major condition that triggered a CA Application Delivery Analysis incident exhibits one clock-hour of normal performance, the management console closes the incident, and Event Manager clears the corresponding event. Afterwards, if the Minor or Major incident condition returns, the management console *opens* a new incident and CA NPC opens a corresponding event.
- If a management console incident remains Open for 24 hours, the management console automatically closes the incident regardless of its condition. If the Minor or Major incident condition still exists, the management console opens a new incident and CA NPC increments the count of the corresponding event. Otherwise, after a synchronization delay of approximately 10 minutes, CA NPC clears the event.

To enable CA NPC to clear an event that is associated with a CA Application Delivery Analysis incident where a server was offline, the management console must report a full hour of No Data for a Minor or Major incident condition for Event Manager to clear the corresponding event. If the management console sees a new Minor or Major condition after a server is brought back online, the management console opens an incident and CA NPC opens a corresponding event.

In CA NPC, if a user closes an event that corresponds to a management console incident, the incident status changes to Acknowledged. After the incident condition exhibits one clock-hour of normal performance, the management console automatically closes the incident.

CA Performance Center (CA PC)

When the management console is registered as a data source with the CA PC, CA Application Delivery Analysis incidents are displayed in the CA PC and managed by the CA Application Delivery Analysis management console. The CA PC does not Acknowledge or Close CA Application Delivery Analysis incidents.

The Performance Events page of the CA PC lists CA Application Delivery Analysis incidents and displays incident counts by server and network. From this page, you can click a CA Application Delivery Analysis incident to drill down to the CA Application Delivery Analysis management console to view incident details and acknowledge the incident:

- If the Minor or Major condition that triggered a CA Application Delivery Analysis incident exhibits one clock-hour of normal performance, the management console closes the incident, and the CA PC clears the corresponding incident. Afterwards, if the Minor or Major incident condition returns, the management console *opens* a new incident and CA PC lists the corresponding incident and increments its incident count.
- If a management console incident remains Open for 24 hours, the management console automatically closes the incident regardless of its condition. If the Minor or Major incident condition still exists, the management console opens a new incident and the CA PC increments its incident count. Otherwise, after a synchronization delay of approximately 10 minutes, CA PC clears the corresponding incident.

To enable CA PC to clear a CA Application Delivery Analysis incident where a server was offline, CA Application Delivery Analysis must report a full hour of No Data for a Minor or Major incident condition. If the management console sees a new Minor or Major condition after a server is brought back online, the management console opens an incident and the CA PC lists the corresponding incident and increments its incident count.

Edit Performance Thresholds

Performance thresholds enable the management console to rate performance and to open incidents. If necessary, you can tighten or loosen the performance thresholds for an application across a group of client networks or all client networks.

Customize the network incident response by customizing performance thresholds by network type.

Use the Performance Thresholds List to manage the list of performance thresholds for each network. In the example below, the DCOM Service Control Manager application has custom thresholds for the client networks in the Austin and San Diego network types. The management console uses the performance thresholds in the Default network type to monitor application performance on any client networks that do not have an assigned network type.

Performance Thresholds List					
Application	TCP Ports	Network Type	Segments	Edit	Delete
DCOM Service Control Manager	135	Default	No		
DCOM Service Control Manager	135	Austin	No		
DCOM Service Control Manager	135	San Diego	No		
Hypertext Transfer Protocol	80	Default	No		
Kerberos Network Authentication Service	88	Default	No		
Lightweight Directory Access Protocol	389	Default	No		
Microsoft Directory Services	445	Default	No		
NetBIOS over IP (MS Windows)	139	Default	No		
Port 1025	1025	Default	No		
Port 1116	1116	Default	No		

Page: 1 Size: 10

Use the Threshold Sets for New Applications list to manage the performance thresholds that are applied to new applications.

You can also create thresholds for a particular network type so that when an application is created or discovered, custom thresholds will be automatically applied to the application on that network. For example, for applications accessed through a VPN, desensitize or disable a performance threshold for network response time. By default, all new applications use the set of performance thresholds established for that network type.

Threshold Sets for New Applications		
Modify threshold sets applied for newly discovered/created applications.		
<input type="button" value="Add Custom by Network Type"/>		
Thresholds Set		
Default		
Internet		
VPN		

Edit Thresholds from the Administration Page

Edit performance thresholds to:

- Set common thresholds for all associated client networks. For example, you can tune the server metrics for an application across all client networks.


In general, a server should treat all of its client requests the same way.

- Set common thresholds for all applications that communicate with a group of client networks. For example, you can tune the network metrics for all the applications that communicate with the client networks that belong to the Austin network type.

If you do not want the management console to include a metric when rating an application, server, or network, you can disable the metric. If you disable a metric, the management console rates an application, server, or network using the available metrics.

To prevent the management console from rating the performance of an application, disable all the Network, Server, and Combined metrics on the application.

Follow these steps:

1. Click the Administration page.
2. Click Policies, Performance Thresholds in the Show Me menu.
3. Scroll to the Performance Thresholds List, find the performance thresholds for the application you want and then click  to edit. The changes you make apply to all client networks with the assigned network type.

Edit Application Thresholds opens.

4. Click Change on the metric you want to edit.
5. Customize the metric thresholds for Minor and Major incidents:
 - Choose a method for measuring a threshold value and then specify a threshold value.
 - Specify the minimum number of observations, or times that the management console must calculate a metric during a 5-minute interval, under Minimum Observations. For example, the management console can calculate Server Connection Time once per TCP transaction, but the management console calculates NRTT each time a client/server pair exchanges a packet. If the management console does not calculate the metric the minimum number of times, the status of the metric is Unrated for the 5-minute interval.

For information about setting threshold properties, click Help.

6. Click Apply.
7. Repeat these steps to edit the Minor and Major threshold values for additional metrics.
8. Click OK.

More information:

[How Performance Metrics Work](#) (see page 141)

Edit Thresholds from the Operations Page

From the Explore button on the Operations page, a management console administrator or network operator can change the threshold value for a particular metric. To change a threshold from the Operations page, isolate the problem by selecting the client network, server, application, and metric. Then, click Explore to edit the threshold for the selected metric.

If the selected network has an assigned network type, edits to the performance thresholds apply to all client networks that belong to the network type. For example, if you change the threshold for NRTT on the Sales application and you have selected a client network that belongs to the London network type, the threshold change applies to all client networks that are assigned to the London network type.

Note For more information about the Explore button, see the *User Guide*.

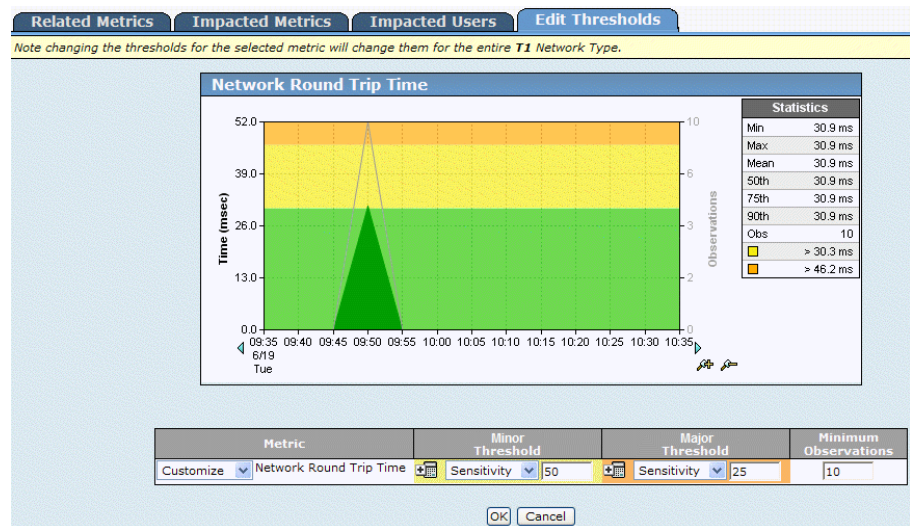
Follow these steps:

1. Click the Operations page.
2. Click Settings.
3. The Setting dialog box opens.
4. Select the client network, server, application, and metric you want.
For information about specifying report settings, click Help.
5. Click Explore.
Metrics Details opens.

- Click the Edit Thresholds tab.

If the selected network has an assigned network type, edits to the performance thresholds apply to all client networks that belong to the network type.

In the following example, the NetQoS LAN network is assigned to the T1 network type. Changes to the LDAP [Client] application performance threshold for NRTT apply to all networks that are assigned to the T1 network type.



- Edit the performance metric to specify the Minor and Major thresholds you want, along with Minimum Observations.
- Click OK.

More information:

[Edit Thresholds from the Administration Page](#) (see page 151)

Add Performance Thresholds

Set tighter or looser performance thresholds for a user-defined application on a particular group of networks. To do this, customize the performance thresholds for the application by network type.

Prerequisite: Assign network types to define a group of networks. Network types allow you to add default performance thresholds to a group of networks.

To automatically monitor a new application with a custom set of thresholds for a group of networks, create performance thresholds for a new application.

Alternatively, you can edit the default threshold settings for an application. The default threshold settings apply to all networks that do not have an assigned network type.

Follow these steps:

1. Click the Administration page.
2. Click Policies, Performance Thresholds in the Show Me menu.
3. Scroll to the Performance Thresholds List.
4. Click Add Custom by Network Type.
Customize Thresholds by Network Type opens.
5. Specify the user-defined application and network type for which you want to define performance thresholds and then click OK.
Edit Application Thresholds opens.
6. Customize the performance thresholds for the application on the network type and click Apply.
7. Click OK.

More information:

[Group Client Networks by Network Type](#) (see page 49)

[Edit Performance Thresholds](#) (see page 150)

[Enable Default Performance Thresholds for a Group of Networks](#) (see page 155)

Enable Default Performance Thresholds for a Group of Networks

Add default performance thresholds to a group of networks to set tighter or looser performance thresholds for newly-discovered applications on that group of networks. The thresholds automatically apply to a new system- or user-defined application.

After you enable default thresholds, they apply to new applications only. Existing applications retain their performance thresholds. Use the Application List, under Data Monitoring, Applications in the Show Me menu, to bulk-edit performance thresholds for existing applications.

Prerequisite: Assign network types to define a group of networks. Network types allow you to add default performance thresholds to a group of networks.

If the application already exists, add performance thresholds to the application for a group of networks.

Follow these steps:

1. Click the Administration page.
2. Click Policies, Performance Thresholds in the Show Me menu.
3. Scroll to the Threshold Sets for New Applications list.
4. Click Add Custom by Network Type.
Customize Thresholds by Network Type opens.
5. Select the network type you want and click OK.
Edit Application Thresholds opens.
6. Customize the performance thresholds for the application by network type.
For information about setting performance thresholds, click Help.
7. Click Apply.
The performance thresholds you specify apply to all newly-discovered applications on a client network that is assigned the network type.
8. Click OK.

More information:

[Edit a User-Defined Application](#) (see page 128)

[Edit a System-Defined Application](#) (see page 114)


Edit Performance Thresholds for WAN-Optimized Network Segments

The management console creates separate applications to report on application performance across the optimized Client, WAN, and Server segments of the network. Customize the performance thresholds on each network segment to make the management console more or less sensitive to variations in application performance. Also, specify performance thresholds for non-optimized application traffic.

Edit Thresholds for Non-Optimized Traffic on an Optimized Application

While you are editing the performance thresholds for each network segment, you can also edit the thresholds for any application traffic that the management console observes which is not optimized.

Follow these steps:

1. Click the Administration page.
2. Click Policies, Performance Thresholds in the Show Me menu.
3. Scroll to the Performance Thresholds List.
4. Click  to edit a WAN-optimized application.

The Segments column identifies WAN-optimized applications. If the management console has observed optimized application traffic by network segment, the status of the Segment column is Yes.

5. Click Thresholds in the third Show Me menu.

If the management console has not calculated response time metrics from non-optimized application traffic, the Thresholds command is not displayed.

Edit Application Thresholds opens.

6. Customize the metric thresholds for Minor (yellow) and Major (orange) incidents:
 - Choose a method for measuring a threshold value and then specify a threshold value.
 - Specify the minimum number of observations, or times that the management console must calculate a metric during a 5-minute interval, under Minimum Observations. If the management console does not calculate the metric the minimum number of times, the status of the metric is Unrated for the 5-minute interval.

For information about setting threshold properties, click Help.

7. Click Apply.
8. Repeat these steps to set threshold values for additional metrics.
9. Click OK.

More information:

[Edit Thresholds from the Administration Page](#) (see page 151)

Edit Thresholds for the Optimized Client Segment

For a WAN-optimized application, edit the performance thresholds for each network segment.

The management console calculates the following metrics from the Client segment:

Network Round Trip Time

Measures the time a packet takes to travel between the server and clients on a network, excluding loss. Application, server, and client processing time are excluded.

Retransmission Delay


Measures the additional delay in the Network Round Trip Time due to retransmissions. The data displayed is an average across all observations, not the actual retransmission time for each transaction.

Transaction Time

Measures the elapsed time from when a client sends the request (packet-level or transaction-level) to when the client receives the last packet in the response.

The management console shows this type of response time data in the Response Time Composition: Average report on the Engineering page. The management console does not open incidents when the Transaction Time threshold is crossed.

Follow these steps:

1. Click the Administration page.
2. Click Policies, Performance Thresholds in the Show Me menu.
3. Scroll to the Performance Thresholds List.
4. Click  to edit a WAN-optimized application. The Segments column identifies WAN-optimized applications. If the management console has observed optimized application traffic by network segment, the status of the Segment column is Yes.
5. Click Client Segment in the third Show Me menu.

If the management console has not calculated response time metrics from the Client network segment, the Client Segment command is not displayed.

Client Segment Thresholds opens.

6. Customize the performance thresholds, Minor (yellow) and Major (orange), for each response time metric:
 - Choose a method for measuring a threshold value and then specify a threshold value.

- Specify the minimum number of observations, or times that the management console must calculate a metric during a 5-minute interval, under Minimum Observations. If the management console does not calculate the metric the minimum number of times, the status of the metric is Unrated for the 5-minute interval.

For information about setting threshold properties, click Help.

7. Click Apply.
8. Repeat these steps to edit the threshold values for additional metrics.
9. Click OK.

More information:

[Edit Thresholds from the Administration Page](#) (see page 151)

Edit Thresholds for the Optimized WAN Segment

For a WAN-optimized application, edit the performance thresholds for each network segment.

The management console calculates the following metrics from the WAN segment:

Network Round Trip Time

Measures the time a packet takes to travel between the server and clients on a network, excluding loss. Application, server, and client processing time are excluded.

Network Connection Time

Measures the time the client takes to confirm the server's connection acknowledgment. Delay is likely to be caused by network latency.


Effective Round Trip Time

Includes Network Round Trip Time plus delays caused by retransmissions. Set a threshold for this metric to monitor performance degradation due to retransmissions.

Retransmission Delay

Measures additional delay in the Network Round Trip Time due to retransmissions. The data displayed is an average across all observations, not the actual retransmission time for each transaction.

Follow these steps:

1. Click the Administration page.
2. Click Policies, Performance Thresholds in the Show Me menu.
3. Scroll to the Performance Thresholds List.
4. Click  to edit a WAN-optimized application. The Segments column identifies WAN-optimized applications. If the management console has observed optimized application traffic by network segment, the status of the Segment column is Yes.
5. Click WAN Segment in the third Show Me menu.

If the management console has not calculated response time metrics from the WAN network segment, the WAN Segment command is not displayed.

WAN Segment Thresholds opens.

6. Customize the metric thresholds for Minor (yellow) and Major (orange) incidents:
 - Choose a method for measuring a threshold value and then specify a threshold value.

- Specify the minimum number of observations, or times that the management console must calculate a metric during a 5-minute interval, under Minimum Observations. If the management console does not calculate the metric the minimum number of times, the status of the metric is Unrated for the 5-minute interval.

For information about setting threshold properties, click Help.

7. Click Apply.
8. Repeat these steps to edit the Minor and Major threshold values for additional metrics.
9. Click OK.

More information:

[Edit Thresholds from the Administration Page](#) (see page 151)

Edit Thresholds for the Optimized Server Segment

For a WAN-optimized application, edit the performance thresholds for each network segment.

The management console calculates the following metrics from the Server segment:

Server Response Time

Measures the time a server takes to start responding to a request made by a client. This value is affected by server speed, application design, and volume of requests.

Server Connection Time

Measures the time a server takes to acknowledge the initial client connection request.


Refused Session Percentage

Measures the percentage of connection requests that were explicitly rejected by the server during the three-way handshake. Part of the Unfulfilled TCP/IP Session Requests report.

Unresponsive Session Percentage

Measures the percentage of sessions where a connection request was sent and the client or server did not respond. Part of Unfulfilled TCP/IP Session Requests report.

Follow these steps:

1. Click the Administration page.
2. Click Policies, Performance Thresholds in the Show Me menu.
3. Scroll to the Performance Thresholds List.
4. Click  to edit a WAN-optimized application. The Segments column identifies WAN-optimized applications. If the management console has observed optimized application traffic by network segment, the status of the Segment column is Yes.
5. Click Server Segment in the third Show Me menu.

If the management console has not calculated response time metrics from the Server network segment, the Server Segment command is not displayed.

Server Segment Thresholds opens.

6. Customize the metric thresholds for Minor (yellow) and Major (orange) incidents:
 - Choose a method for measuring a threshold value and then specify a threshold value.
 - Specify the minimum number of observations, or times that the management console must calculate a metric during a 5-minute interval, under Minimum Observations. If the management console does not calculate the metric the minimum number of times, the status of the metric is Unrated for the 5-minute interval.

For information about setting threshold properties, click Help.

7. Click Apply.
8. Repeat these steps to edit the threshold values for additional metrics.
9. Click OK.

More information:

[Edit Thresholds from the Administration Page](#) (see page 151)

Chapter 7: Managing Incident Responses

This section contains the following topics:

[How Incident Responses Work](#) (see page 165)

[Add an Incident Response](#) (see page 175)

[Edit an Incident Response](#) (see page 176)

[Delete an Incident Response](#) (see page 177)

[Add an Action to a Network or Server Incident Response](#) (see page 177)

[Edit a Responsive Action](#) (see page 178)

[Delete a Responsive Action](#) (see page 178)

[Assign an Incident Response](#) (see page 179)

[Troubleshoot Incident Responses](#) (see page 181)

[Manage Incidents Using Web Service Methods](#) (see page 181)

How Incident Responses Work

To help you troubleshoot a problem at the time that it occurs, and reduce the mean time to repair, assign incident responses to your business-critical applications, servers, and networks. Incident responses:

- Notify your team about a performance degradation.
- Actively investigate a problem to gather additional information that can help you identify the root cause of a performance degradation.

By default, the management console does not automatically launch an incident response.

Note that you can also manually launch an investigation from the Incidents page to further troubleshoot an incident. For more information, see the *User Guide*.

How an Incident Response is Launched

When the management console sees a threshold violation for a network or server metric, the management console automatically opens a network or server incident. Available from the Incidents page, an *incident* creates a record of information about a performance problem on a server or network, including the affected applications.

When the management console opens an incident for a:

- Server, the management console evaluates the server incident and launches a set of actions on the affected server and to the poor performing applications running on that server.
- Network, the management console evaluates the network incident and launches a set of actions on the network and to the poor performing applications running across the network.

Although the management console rates the performance of an application using Combined metrics, the management console does not create application incidents. However, the management console lets you define application incident responses. An *application incident response* is an application response to a network or server incident. For example, if you configure an application incident response for the Exchange application, the management console launches the incident response when a:

- Network incident is created by clients accessing the Exchange application
- Server incident is created by a server that hosts the application

The management console does not launch an application incident response when the threshold for a combined metric, such as Data Transfer Time, is crossed.

By default, the management console does not launch a notification or investigation in response to a network or server incident. Edit the Default incident responses to add one or more actions, and create additional incident responses as needed.

Network Incident Responses

A network incident response is launched in response to a network incident. The following responses are available to a network incident:

- [Email Notification](#) (see page 168)
- [SNMP Trap Notification](#) (see page 169)
- [Trace Route Investigation](#) (see page 174)

Server Incident Responses

A server incident response is launched in response to a server incident. The following responses are available to a server incident:

- [Email Notification](#) (see page 168)
- [SNMP Trap Notification](#) (see page 169)
- [Ping Response Time Investigation](#) (see page 173)
- [Performance via SNMP Investigation](#) (see page 172)
- [Packet Capture Investigation](#) (see page 171)

Application Incident Responses

An application incident response is launched in response to a network or server incident. The following application responses are available:

- [Email Notification](#) (see page 168)
- [SNMP Trap Notification](#) (see page 169)
- [Application Connection Time Investigation](#) (see page 170)

Note that the management console does not launch this investigation in response to an application that is rated as Unavailable.

- [Packet Capture Investigation](#) (see page 171)

More information:

[How Server Incidents for Application Availability Work](#) (see page 203)

Email Notifications

Email notifications update someone about the status of the affected applications, servers, or networks.

The management console sends email notifications to the SMTP server that is specified by CA PC. Otherwise, the management console sends email notifications to the SMTP server that is specified by the CA Application Delivery Analysis console settings. The CA NPC email specifications are used to send scheduled or ad hoc reports from the CA NPC. The SMTP server configured in the CA NPC is not used by the management console to send email notifications.

To include more than one application, server, or network in the same email, assign the same incident response to more than application, server, or network type.

If an assigned server or network meets the incident Duration and Severity criteria, the management console includes this information in a status update email that the management console sends every two hours.

Edit the Duration and Criteria thresholds to send email notifications about:

- Minor performance degradations that occur occasionally for short periods of time. Something might be happening that you should monitor.
- Major performance degradations that last for 10 minutes or an Unavailable condition that lasts for 5 minutes. You should investigate these types of situations right away.

More information:

[Manage Console Settings](#) (see page 223)

[How Application Availability Reporting Works](#) (see page 203)

SNMP Trap Notifications

Use a SNMP trap notification to update a SNMP Manager about the Open or Closed incident status of the affected applications, servers, or networks.

You can assign a SNMP trap notification to any incident response. The management console sends SNMPv2 trap notifications from the management console computer.

Unlike email notifications, the management console sends a SNMP trap when the management console:

- Opens and closes a server or network incident
- (Optional) A severity change occurs on a server or network incident, for example, from Minor to Major.

To include more than one application, server, or network in the same SNMP trap, assign the same incident response to more than application, server, or network type. If more than one application, server, or network is affected, the SNMP trap includes a URL with the details.

By default, the management console sends SNMP traps using the SNMPv2 community named *SuperAgent*. The default SNMP profile only applies SNMP trap notifications.

To enable a SNMP manager to understand a SNMP trap notification from the management console, compile the management console MIB into the trap receiver on the SNMP manager. The compiling method varies depending on the SNMP manager.

The management console MIB is not included with the management console distribution. Download the CA Application Delivery Analysis MIB from the CA Support website at <http://support.ca.com>.

CA Application Delivery Analysis trap notifications use the following abbreviations for performance metrics.

Abbreviation	Metric
ERTT	Effective Round Trip Time
NCT	Network Connection Time
NRTT	Network Round Trip Time
RS	Refused session (percentage)
RTNS	Retransmission delay
SCT	Server Connection Time
SRT	Server Response Time

Abbreviation	Metric
US	Unresponsive session (percentage)

More information:

[Interpret SNMP Traps](#) (see page 187)

Application Connection Time Investigations

Use an application connection time investigation to gather information about the time it takes to connect to a TCP/IP application port. This includes time for the server to respond with a connection acknowledgment.

You can assign an application connection time investigation to an application incident response. If the server that hosts the application is monitored by a CA Standard Monitor, the management console launches this investigation from the monitoring device. Otherwise, the management console launches the investigation from the management console.

Note that the management console does not launch this investigation in response to an Unavailable application.

A management console administrator can also launch or schedule this investigation from the Incidents page. For more information, see the *User Guide*.

An application connection time investigation produces a report about the number of successful and unsuccessful attempts and connection time.

More information:

[How Server Incidents for Application Availability Work](#) (see page 203)

Packet Capture Investigations

Use a packet capture investigation to perform a filtered capture of the particular server, application port, and network experiencing a problem.

You can assign a Packet Capture investigation to an application in response to a server incident, or to a server and to the poor performing applications running on that server. The management console automatically takes the packet capture from the appropriate monitoring device. When the packet capture is taken by:

- A CA Standard Monitor, the capture file is copied from the monitoring device to the user's computer on TCP-8080. Make sure the user's computer has access to the monitoring device.
- A CA Multi-Port Monitor, the capture file is not copied to the user's computer. The user views the capture file on the monitoring device.

The following monitoring devices do not take a packet capture investigation for CA Application Delivery Analysis:

- Cisco WAE device
- Cisco NAM device

A management console administrator can also launch or schedule this investigation from the Incidents page. For more information, see the *User Guide*.

A packet capture investigation report includes information about the server, application, network, and a link to view the packet capture results.

Performance via SNMP Investigations

Use a performance via SNMP investigation to SNMP poll a server for its performance information, such as memory and CPU utilization.

Assign a performance via SNMP investigation to a server incident response. If the server is monitored by a CA Standard Monitor, the management console launches this investigation from the monitoring device. Otherwise, the management console launches the investigation from the management console.

From the Incidents page, a management console administrator can also launch or schedule this investigation on a server or router. To SNMP poll a router for performance information, the management console administrator must add the router to the management console as a network device.

- If you have not assigned a valid SNMP profile to the server or network device, the management console attempts to discover a valid SNMP profile. Alternatively, you can use another CA product to SNMP poll your servers and routers for performance information.
- To enable the management console to SNMP poll a server or network device, you must add a valid SNMP profile to the management console.

The management console uses the following criteria to report the overall processor and process CPU utilization values:

- Individual processor percentages are polled using HrProcessorTable (1.3.6.1.2.1.25.3.3.1.2) maintained by the host over the last minute (minute before first poll). This is a snapshot.
- The individual process lists are polled using HrSWRunPerfTable (1.3.6.1.2.1.25.5.1.1.1 and .2, cpu and mem). The CPU value is a raw time; therefore, it is polled twice and subtracted. The percentage is derived from the total CPU time used divided by available CPU time. Memory percentage is the last poll divided by available memory.

The investigation polls a server twice during an investigation--once when the incident is opened and again 5 minutes later, for the following MIBs:

- System mibs (1.3.6.1.2.1.1.*.0)
- Interface mibs (1.3.6.1.2.1.2.2.1)
- Cisco CPU mibs (1.3.6.1.4.1.9.2.1.*)
- Host resource processor (1.3.6.1.2.1.25.3.3.1)
- Host resource storage (1.3.6.1.2.1.25.2.3.1)
- IP address table (1.3.6.1.2.1.4.20.1)
- Host resource running software mib (1.3.6.1.2.1.25.4.2.1)
- Host resource running software performance mib (1.3.6.1.2.1.25.5.1.1)

- Host resource memory (1.3.6.1.2.1.25.2.2.0)

When launched or scheduled on a:

- Server, a performance via SNMP investigation produces a report about server performance and interface statistics.
- Network device, a performance via SNMP investigation produces a report about device performance and interface statistics.

Ping Response Time Investigations

Use a ping response time investigation to verify that a server can respond to a ping request and measure the round-trip time to receive the response.

Assign a ping response time investigation to a server incident response. If the server is monitored by a CA Standard Monitor, the management console launches this investigation from the monitoring device. Otherwise, the management console launches the investigation from the management console.

A user can also launch or schedule this investigation from the Incidents page. For more information, see the *User Guide*.

A ping response time investigation produces a report about packet round trip time and response time.

Trace Route Investigations

Use a trace route investigation to record the path and each hop between the monitoring device and end-points to monitor latency and routing issues, and optionally, SNMP poll each router for its performance information.

Assign a trace route investigation to a network incident response. If the server that hosts the application is monitored by a CA Standard Monitor, the management console launches this investigation from the monitoring device. Otherwise, the management console launches the investigation from the management console.

During a network incident, the management console will begin a Trace Route investigation from the monitoring device and attempt to reach a device on the network subnet with the incident. If the subnet is a /24 or is more specific, an IP address in the subnet is used as a target; the management console must have collected response time data for this address. If the subnet is less specific than a /24, the management console picks an address in the subnet range using another method. If within the Trace Route action properties, the option "Investigate Routers via SNMP" is enabled, then the trace route is performed first, and afterwards, the list of network devices in the route is used for SNMP investigation. If the management console discovers a valid SNMP profile, the management console queries the device for its performance information, which is included in the investigation results.

When configured to run a TCP trace route, the TCP trace route uses the TCP port of the applications being monitored for outbound traffic and ICMP TTL Expired messages for the return to isolate the router along the path that killed the outbound packet.

A management console administrator can also launch or schedule this investigation from the Incidents page.

A trace route investigation produces a report about paths, hops, delay, and usage.

Optionally, you can configure a trace route investigation to launch a Performance via SNMP investigation of device and performance interface statistics for each device in the path.

To SNMP poll each hop along the path, you do not need to add network devices to the management console, however, to SNMP poll a server or network device, you must define a valid SNMP profile. A performance via SNMP investigation produces a report about device performance and interface statistics.

To SNMP poll each hop along the path, if you have not defined a network device in the management console with a valid SNMP profile, the management console attempts to discover a valid SNMP profile on each network device. Alternatively, you can use another CA product to SNMP poll your devices for performance information.

Add an Incident Response

To enable the management console to send a notification or launch an investigation in response to a network or server incident:

1. Add a network, server, or application incident response.
2. Add one or more actions to the incident response.
3. Assign the incident response to a particular network, server, or application.

Alternatively, edit the default network, server, and application incident response to add a responsive action. By default, the management console does not launch a notification or investigation in response to a network or server incident.

For the management console to launch a notification, both the incident Duration and Severity criteria must be met. Use the Duration and Severity criteria to launch an incident response for:

- Minor performance degradations that occur occasionally for short periods of time. Something might be happening that you should monitor.
- Major performance degradations that last for 1 hour or an Unavailable condition that lasts for 10 minutes. You should investigate these types of situations right away.

Follow these steps:

1. Click the Administration page.
2. Click Policies, Incident Responses in the Show Me menu.
3. Under the Show Me menu, click to add an action.

Add Network Response

Notifies the owner of the client networks for a remote site, or runs a trace route investigation.

Add Server Response

Notifies the owner of the server, or runs a server-based investigation, such as a packet capture, ping response time, or performance via SNMP investigation.

Add Application Response

Notifies the owner of the application, or runs an application connection time investigation.

4. Click Apply.

You are ready to [add an action to the incident response](#) (see page 177).

More information:

[Managing Application Availability](#) (see page 201)


[Assign an Incident Response](#) (see page 179)

[Manage Monitoring Device Incidents](#) (see page 253)

Edit an Incident Response

Edit an incident response to rename it. While you are renaming an incident response, you can also modify its responsive actions, for example, to edit an action or add an action.

Follow these steps to rename an incident response:

1. Click the Administration page.
2. Click Policies, Incident Responses in the Show Me menu.
3. Browse the incident lists and click  to edit a network, server, or application incident response.

Verify at least one responsive action is assigned to each incident response.

4. Click Edit Incident Response in the third Show Me menu.
5. Type a new name in Incident Response Name and click OK.

More information:

[Add an Action to a Network or Server Incident Response](#) (see page 177)

[Edit a Responsive Action](#) (see page 178)


[Delete a Responsive Action](#) (see page 178)

Delete an Incident Response

Deleting an incident response deletes the incident response and its responsive actions. If the incident response is assigned, the management console reassigns the default incident response to any affected applications, servers, or networks.

You cannot delete the default incident response for a network, server, or application.


Follow these steps:

1. Click the Administration page.
2. Click Policies, Incident Responses in the Show Me menu.
3. Browse the incident response lists and click  to delete a network, server, or application incident response. Note that you cannot delete the Default incident responses.
4. In the Delete Confirmation, click Continue with Delete to delete the incident response.

Add an Action to a Network or Server Incident Response

Add an action to launch a notification or investigation in response to a network or server incident.

Follow these steps:



1. Click the Administration page.
2. Click Policies, Incident Responses in the Show Me menu.
3. Browse the incident response lists and click  to edit a network, server, or application incident response.
4. Click Edit Actions in the third Show Me menu.
5. Click Add Action under the Show Me menu.
Action Types opens.
6. Select an action and click Next.
Action Properties opens.
7. Specify the responsive action settings and click OK. For more information, click Help.

In Incident Response Actions, the new action is displayed.

Edit a Responsive Action

Edit the default network, server, and application incident responses to add one or more responsive actions, and create additional incident responses as needed. By default, the management console does not launch a notification or investigation in response to a network or server incident.

Follow these steps:

1. Click the Administration page.
2. Click Policies, Incident Responses in the Show Me menu.
3. Browse the incident response lists and click  to edit a network, server, or application incident response.
4. Click Edit Actions in the third Show Me menu.
Incident Response Actions opens.
5. Click  to edit an action.
Action Properties opens.
6. Specify the responsive action settings and click OK. For more information, click Help.



More information:

[Manage Monitoring Device Incidents](#) (see page 253)

Delete a Responsive Action

If you no longer want the management console to launch a particular responsive action, delete the action from its network, server, or application incident response.

Follow these steps:

1. Click the Administration page.
2. Click Policies, Incident Responses in the Show Me menu.
3. Browse the incident response lists and click  to edit a network, server, or application incident response.
4. Click Edit Actions in the third Show Me menu.
Incident Response Actions opens.
5. Click  to delete an action.
6. Click Continue with Delete at the prompt to delete the responsive action.

Assign an Incident Response

To enable the management console to respond to a network or server incident, assign an incident response with at least one responsive action to an application, server, or network type.

Note that the default incident responses do not include a responsive action, but you can add them.

Assign an incident response to	To launch an action on
An application	An application with a Minor (yellow) or Major (orange) performance degradation on a server or network.
A server	A server with a Minor (yellow) or Major (orange) performance degradation.
A network type	A network with a Minor (yellow) or Major (orange) performance degradation. Note that when you assign a incident response to a network type, the incident response is assigned to all networks with the corresponding network type.

More information:

[Group Client Networks by Network Type](#) (see page 49)

[Edit a Responsive Action](#) (see page 178)

Assign an Incident Response to an Application

Assign an incident response to an application to enable the management console to investigate the application and notify your team about a network or server incident that affects the application.



Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Applications in the Show Me menu.
3. Scroll to the Application List, select the applications you want, and click Edit.
Application Properties opens.
4. Click Incident Response to specify the incident response you want and click OK.

Assign an Incident Response to a Server

Assign an incident response to a server to enable the management console to investigate the server and notify your team about a corresponding server incident. The management console assigns a default incident response to all servers, however, CA Application Delivery Analysis is a passive monitoring solution and by default, does not take any action.

Follow these steps:


1. Click the Administration page.
2. Click Data Monitoring, Applications in the Show Me menu.
3. Scroll to the Application List and click  to edit a server.
Server Properties opens.
(Optional) To assign the same incident response to more than one server, select the appropriate servers, and click  to edit all selected applications.
4. In the Server Properties, click Incident Response to assign the server incident response you want and click OK.

Assign an Incident Response to a Network Type

Assign an incident response to a network type to enable the management console to investigate a network and notify your team about a corresponding network incident.

Note that when you assign an incident response to a network type, the same incident response applies to all networks with the corresponding network type.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Network Types in the Show Me menu.
Network Types opens.
3. Click  to edit a network type.
Network Type Properties opens.
4. Click Incident Response to specify the incident response you want and click OK.

More information:

[Assign a Network Type to a Client Network](#) (see page 54)

Troubleshoot Incident Responses

If the management console does not launch the responsive action you expect, make sure:

- The appropriate incident response is assigned to the application, server or network type. For a client network, make sure the appropriate network type is assigned to the client network.
- The responsive action that is assigned to the incident response is configured properly. In general, verify the network or server incident created by the management console exceeds the Minimum Severity and Duration criteria for the incident response. Review the responsive action properties and see the Help for more information.

Manage Incidents Using Web Service Methods

Use Web services to query for incident data beyond what gets sent when you add a SNMP trap action to an incident response. Use the Send Status Updates option with incident response traps.

Object Identifier Specifications

- Incident Id .1.3.6.1.4.1.4498.2.20.1.1.1
- Server Name .1.3.6.1.4.1.4498.2.20.1.2.1
- Server IP .1.3.6.1.4.1.4498.2.20.1.3.1
- Application Name .1.3.6.1.4.1.4498.2.20.1.4.1
- Network Name .1.3.6.1.4.1.4498.2.20.1.5.1
- Metric Name .1.3.6.1.4.1.4498.2.20.1.6.1
- Incident Time .1.3.6.1.4.1.4498.2.20.1.7.1
- Severity .1.3.6.1.4.1.4498.2.20.1.8.1
- Impact % .1.3.6.1.4.1.4498.2.20.1.9.1
- Duration .1.3.6.1.4.1.4498.2.20.1.10.1
- Incident URL .1.3.6.1.4.1.4498.2.20.1.11.1
- Response Type .1.3.6.1.4.1.4498.2.20.1.12.1 Server | Network | Application
- Incident State Type .1.3.6.1.4.1.4498.2.20.1.13.1 Open | Update | Close
- Web Service IP .1.3.6.1.4.1.4498.2.20.1.14.1

Server Name and IP

- For all server incidents, the Server Name and Server IP fields contain the specific name and IP address of the associated server.
- For network incidents with only one related server, these fields contain the specific name and IP address of the associated server.
- For network incidents with multiple related servers, these fields contain [n+ Servers] or [n+ Addresses] to indicate the number of Severity state servers, n or more.
- For network incidents closing because of data inactivity (white status), these fields are [None].

Application Name

- For server or network incidents with only one related application, this field contains the specific name for that application.
- For application incident responses to server or network incidents, this field contains the specific name for that application.
- For server or network incidents with more than one related application, this field contains [n Applications] where n is the number of severity state applications.
- For server or network incidents closing because of data inactivity (white status), this field is [None].
- For server incidents caused by server unavailability, this field is [All].

Network Name

- For all network incidents, this field contains the specific name and subnet of the associated network.
- For server incidents with only one related network, this field contains the specific name and subnet of the associated network.
- For server incidents with multiple related networks, this field contains [n+ Networks] to indicate the number of Severity state networks, n or more.
- For server incidents closing because of data inactivity (white status), this field is [None].

Metric Name

- For incidents caused by a metric crossing a threshold, this field is a comma delimited list of abbreviations.
 - NRTT Network Round Trip Time
 - RTNS Retransmission Delay
 - NCT Network Connection Time
 - ERTT Effective Round Trip Time
 - SRT Server Response Time
 - SCT Server Connection Time
 - RS Refused Session (percentage)
 - US Unresponsive Session (percentage)
- For incidents caused by unavailability, this field is Availability.
- For incidents closing because of data inactivity (white status), this field is Insufficient data.

Web Service Specifications

Use one of the Web service interface methods to obtain data. These Web services support XML output.

Access Web services here:

<http://WebServiceIP/SuperAgentWebService/PublishedService.asmx>

Input

Retrieve server lists using `FetchServersByIncident` or `FetchServersByIncidentSeverityDuration`.

Input

- Incident ID
- Minimum Severity (optional)
- Minimum Duration (optional)

Output

- `server_id`, 32 bit unsigned identifier
- `server_desc`, maximum 50 char string name
- `address`, 32 bit unsigned IP address
- `None`, floating point percentage of time during the incident in which no data was collected
- `Unrated`, floating point percentage of time during the incident in which no threshold existed (had not yet been calculated or was turned off)
- `Normal`, floating point percentage of time during the incident in which normal (below threshold) data was collected
- `Degraded`, floating point percentage of time during the incident in which degraded (but not excessive) data was collected
- `Excessive`, floating point percentage of time during the incident in which excessive data was collected
- `Unavailable`, floating point percentage of time during the incident for which the server or application was unavailable (server incidents only)

Input

Retrieve application lists `FetchApplicationsByIncident` or `FetchApplicationsByIncidentSeverityDuration`.

Input

- Incident ID
- Minimum Severity (optional)
- Minimum Duration (optional)

Output

- `app_id`, 32 bit unsigned identifier
- `applications_desc`, maximum 50 char string name
- `port_beg`, 16 bit unsigned first port in the range
- `port_end`, 16 bit unsigned last port in the range
- `None`, floating point percentage of time during the incident in which no data was collected
- `Unrated`, floating point percentage of time during the incident in which no threshold existed (had not yet been calculated or was turned off)
- `Normal`, floating point percentage of time during the incident in which normal (below threshold) data was collected
- `Degraded`, floating point percentage of time during the incident in which degraded (but not excessive) data was collected
- `Excessive`, floating point percentage of time during the incident in which excessive data was collected
- `Unavailable`, floating point percentage of time during the incident for which the server or application was unavailable (server incidents only)

Input

Retrieve network lists `FetchNetworksByIncident` or `FetchNetworksByIncidentSeverityDuration`.

Input

- Incident ID
- Minimum Severity (optional)
- Minimum Duration (optional)

Output

- `client_id`, 32 bit unsigned identifier for the network definition
- `client_address`, 32 bit unsigned address component of the network subnet
- `client_mask`, 32 bit unsigned mask component of the network subnet (expanded CIDR)
- `client_desc`, maximum 50 char string name
- `subnet`, x.x.x.x/m formatted subnet specification
- None, floating point percentage of time during the incident in which no data was collected
- Unrated, floating point percentage of time during the incident in which no threshold existed (had not yet been calculated or was turned off)
- Normal, floating point percentage of time during the incident in which normal (below threshold) data was collected
- Degraded, floating point percentage of time during the incident in which degraded (but not excessive) data was collected
- Excessive, floating point percentage of time during the incident in which excessive data was collected
- Unavailable, floating point percentage of time during the incident for which the server or application was unavailable (server incidents only)

Note: The severity percentages from incident web services are consistent with Operations and Incidents view exports.

Send Status Updates Specifications

An application, server, and network incident response trap option sends incremental severity updates. All severity updates are sent regardless of duration. If you turn off this option, only open and close incident traps are sent on sufficient severity/duration.

Application Incident Response Trap Specifications

Application incident response traps are sent as network and server incident responses specific only to the application. Normally, network incidents contain aggregate application and server data. Similarly, a server incident contains aggregate application and network data. With application incident responses to network and server incidents, every network/application and server/application combination over threshold will alarm and contain aggregate server and network data, respectively.

Interpret SNMP Traps

The following table lists lines from a trap with descriptions of each line.

Trap	Description
1.3.6.1.4.1.4498.2.20.1.1.1 netQoSIncident7Number 17841	Incident identifier.
1.3.6.1.4.1.4498.2.20.1.2.1 netQoSIncident7Server dc1.netqos.local	Name of the server involved.
1.3.6.1.4.1.4498.2.20.1.3.1 netQoSIncident7ServerAddress 192.168.0.6	Address of the server involved.
1.3.6.1.4.1.4498.2.20.1.4.1 netQoSIncident7Application Lightweight Directory Access Protocol	Application involved.
1.3.6.1.4.1.4498.2.20.1.5.1 netQoSIncident7ClientRegion SuperAgent LAN - 192.168.245.0/24	Network involved.
1.3.6.1.4.1.4498.2.20.1.6.1 netQoSIncident7Regards NRTT,ERTT	Metrics involved.
1.3.6.1.4.1.4498.2.20.1.7.1 netQoSIncident7Time 02/20/2007 20:40 GMT-4	End stamp of the event.
1.3.6.1.4.1.4498.2.20.1.8.1 netQoSIncident7Severity Excessive	Threshold violated and the severity classification.

Trap	Description
1.3.6.1.4.1.4498.2.20.1.9.1 netQoSIncident7Impact 91.5%	Weighed percentage of observations this affects. The impact value is the impact of the netQoSIncident7Severity for the previous netQoSIncident7Duration minutes. For networks, it is the peak percent of server observations per app/metric pair. For servers, it is the peak percent of network observations per app/metric pair.
1.3.6.1.4.1.4498.2.20.1.10.1 netQoSIncident7Duration 10.0 Minutes	Duration of the event.
1.3.6.1.4.1.4498.2.20.1.11.1 netQoSIncident7URL http://192.168.100.131/SuperAgent/Investigator/Incidents/IncidentsViewFocus.aspx?Nav=13,0,0&Stack=T M N A S&I=1017841	URL link to the UI for the incident.
1.3.6.1.4.1.4498.2.20.1.12.1 netQoSIncident7ResponseType Network	Type of incident response that generated this trap.
1.3.6.1.4.1.4498.2.20.1.13.1 netQoSIncident7State Open	The Open, Update, or Close state of the trap.
1.3.6.1.4.1.4498.2.20.1.14.1 netQoSIncident7WebServiceIP 192.168.100.131	The IP address of the Console where Web services for further details can be found.

Chapter 8: Managing Application Performance OLAs

This section contains the following topics:

[How Performance OLAs Work](#) (see page 189)

[Establish Operational Levels from Historical Data](#) (see page 194)

[Create an Application Performance OLA for a Group of Networks](#) (see page 196)

[Edit an Application Performance OLA](#) (see page 198)

[Delete an Application Performance OLA](#) (see page 199)

How Performance OLAs Work

A *performance operational level agreement* (performance OLA) lets you evaluate compliance with application performance goals for a remote site. By default, the management console does not define operational levels for application performance.

Performance OLAs enhance reporting by tracking the behavior of the worst performing IPv4-based transactions over time. This tracking indicates where and when performance degradation is most serious. This tracking also enables users to understand how performance varies from the mean data points reported in the management console.

By default, the management console does not report performance OLAs. You can create OLAs for a user-defined application, but not for a system-defined application.

Because of the way the management console collects OLA data and measures OLA compliance, it is recommended that you set up OLAs to measure compliance at each remote location rather than across all locations. To establish OLAs at each remote location, use network types.

(Optional) To set a performance OLA for an application and apply the OLA to all the servers monitored by the management console, create a user-defined application and assign it to a domain. The management console automatically keeps the application server assignments up-to-date as your server subnets change.

More information:

[Group Client Networks by Network Type](#) (see page 49)

How Performance OLA Reporting Works

A performance OLA shows how well a user-defined application is performing by counting, on an hourly basis, the percentage of IPv4-based transactions that are faster than a given threshold. An example of a performance OLA would be that 90 percent of server response times must be under 20 milliseconds. The result of a performance OLA is displayed in a report which indicates, on an hourly basis, whether the application meets the operational level agreement.

The management console measures hourly operational level compliance based on the threshold you specify. You can specify a threshold with resolution of up to 3 decimal places, such as 93.999. It can take up to an hour to report an OLA compliance violation.

The management console does not collect OLA data in the same manner as 5-minute data. For performance OLAs, the monitoring device:

- Compares every transaction to see whether it meets the OLA thresholds.
- Records how many transactions pass and fail every hour.


An analogy would be, instead of recording the average speed of cars on an interstate every five minutes, the management console marks whether each is traveling above 40mph and reports the successes and failures each hour.

By comparison, the management console uses thresholds to monitor application performance based on the average response time of a particular metric during a 5-minute interval.

How Performance OLA Thresholds Work

Because performance OLAs give additional insight into the slowest performing transactions, it is important to understand what a typical time slice of performance looks like. Most people are familiar with a standard distribution or bell curve, and when talking in terms of percentiles these usually come to mind. However, TCP transactions do not follow a normal distribution. Instead, there is minimum latency that is due to distance for Network Round Trip Time and I/O in Server Response Time (SRT). Ideally, you want to have as many transactions as close to this minimum as possible.

The following example shows an idealized version of what the management console sees over a given time period for Network Round Trip Time. When performance degrades, the entire curve could shift to the right, or the tail in the curve could extend. This situation is one where OLAs are useful.



OLAs enable the user to specify a threshold for the 90th and 98th percentiles. Note that the user can specify any percentile. Reports show what percentile of those values actually reached the threshold. Adjusting the thresholds lets you monitor tail behavior and set OLAs in a goal-oriented manner.

How Operational Level Metrics Work

A performance OLA uses the following key metrics to measure operational level compliance. You can choose to monitor some or all of these operational levels in your OLA:

Operational Level Metric	Details
Network Round Trip Time	<ul style="list-style-type: none">■ Configure by network type, especially for WAN links.■ Verify that the networks with similar latency are grouped together.■ Monitor business critical applications, especially for WAN links.■ Select applications that represent the largest (or most constant) amount of traffic to each site to get more observations and statistical significance.
Server Response Time	<ul style="list-style-type: none">■ Do not separate by network type.■ Servers should treat requests independently no matter which network they came from.■ Monitor the most influential tier in a multi-tiered application and keep in mind that front end server response time will include back-end server requests.■ Monitor servers for business critical applications
Total Transaction Time	<ul style="list-style-type: none">■ Best indicator of overall end user experience because it is dependent on Server Response Time, Network Round Trip Time, and Data Transfer Time.■ Configure on a per network type basis because it depends on network round trip time.■ Monitor business critical applications

Performance OLA Tips and Tricks

We recommend creating an application performance OLA for user-defined applications that represent the largest (or most constant) amount of traffic to each site to get more observations and statistical significance. You cannot set operational levels for system-defined applications.

When setting up a performance OLA:

- Separate data center metrics from network metrics. Not every metric is meaningful for every network. For example, on back-end applications, you might not want to see Network Round Trip time tracking.
- It is appropriate to use longer timelines when deciding OLA threshold values to eliminate transient spikes or dips. You can also exclude a particular metric from the OLA.
- Define up to two thresholds to measure application compliance with an operational level, for example, to set internal and external operational levels. In general, define an operational level at 90 percent or higher so that the management console can offer additional insight into the variability of performance. Alternatively, you are not required to set performance threshold levels for an OLA metric.

With the correct value for each threshold level, the management console will have historical trend data on the 5-minute averages and on the worst performing (highest percentile) transactions that the users in each remote location are experiencing.

More information:


[Edit an Application Performance OLA](#) (see page 198)

Establish Operational Levels from Historical Data

Use the historical report data as a starting point for configuring a new operational level agreement. We recommend that you wait a month to allow the management console to collect enough data to determine normal behavior for the application.

Follow these steps:

1. Click the Engineering page.
2. Click Settings on the Show Me menu.
The Settings dialog box opens.
3. Select an application.
4. Select a server.
5. Select a network.
6. Set the time frame to display a monthly report and click OK.
7. Find the operational level metric you want in the Components Report, such as the Network Round Trip Time report which, in the example below, contains response time data for the Los Angeles network and all the servers that host the Exchange application.



8. In Statistics, make a note of the 90th and Max values. In the example below, the 90th percentile is 12.4 ms and the maximum value is 41.7 ms.

Statistics	
Min	5.52 ms
Max	41.7 ms
Mean	15.6 ms
50th	14.4 ms
75th	17.4 ms
90th	12.4 ms
Obs	172M

Note When reporting on a monthly time frame, the management console aggregates reporting statistics into 6-hour increments. For more information, see the *User Guide*.

9. Repeat these steps to gather the 90th and Max percentile for each of the operational level metrics, including Network Round Trip Time and Server Response Time. For Transaction Time, sum and average the 90th and Max percentile of Network Round Trip Time, Server Response Time, Retransmission Delay, and Data Transfer Time.
10. To determine the initial operational levels, double the value of the historic monthly 90th and 98th percentiles. Based on the Statistics example above, doubling the 90th and 98th percentiles would equate to a Network Round Trip Time 90th percentile of 25 milliseconds and a 98th percentile of 81 milliseconds.
11. Monitor the OLA as it runs over the first reporting period; do not be discouraged if it does not meet OLA at first. Wait for a full reporting period to complete before tuning. Because the management console measures operational level compliance on an hourly basis, there can be up to an hour lag before an OLA shows a compliance violation.
12. Tune the thresholds after a reporting period completes to set thresholds that are achievable. The results enable you to make a more educated assessment of the thresholds.

More information:

[Create an Application Performance OLA for a Group of Networks](#) (see page 196)

Create an Application Performance OLA for a Group of Networks

Create a performance OLA to define the expected levels of service for a user-defined application across a group of client networks with similar latency characteristics.

The management console creates a default OLA for each user-defined application to measure performance across all networks that do not have an assigned network type. In general, this is not recommended for monitoring Network Round Trip Time and Total Transaction Time because of differences in network latency between remote locations.

When you update the threshold values in a default OLA, your changes are dynamically applied to any new OLAs for that application that are configured to use the default OLA's value. For example, if the default OLA for the POP3 application requires at least 90% of the observations for Network Round Trip Time during a 1-hour interval to be at least 5ms, and you create an OLA for the POP3 application in the Austin network type, the default percentage of observations is 90% and the Network Round Trip Time threshold is 5 ms. If you change POP3's default OLA threshold for Network Round Trip Time, the management console dynamically updates the Network Round Trip Time threshold for the POP3 Austin's OLA to use the new default value.

When creating an OLA, you can customize a threshold value to override the dynamic default value. If you change the number of thresholds or minimum observations for an OLA metric without customizing the threshold values, the change applies to the default OLA.

Before you create an OLA, it is recommended that you monitor an application for a month to determine its normal performance.

Follow these steps:

1. Click the Administration page.
2. Click Policies, Performance OLAs in the Show Me menu.
3. Browse the list of OLAs in the Performance OLA List. If there are no applications in the list, no user-defined applications exist.

By default, the management console creates a Default OLA for each user-defined application that applies to networks without an assigned network type.

4. Click Add Custom by Network Type to create an application OLA for a group of client networks.

Customize Thresholds by Network Type opens.

5. Select the application and network type that corresponds to the appropriate remote location and click OK.

Edit OLA Thresholds opens.

6. Edit the threshold settings for each OLA metric and click OK.

For information about editing OLA thresholds, click Help.

7. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

8. Wait a full reporting cycle, up to one hour, to see how the threshold changes compare to the actual response times observed by the management console.
 - a. On the Management page, in the Show Me menu, click Performance Detail OLA.
 - b. If necessary, adjust the values for a operational level threshold percentage or increase the value of the threshold level so that the application meets the operational level agreement for the remote site.

More information:

[Edit an Application Performance OLA](#) (see page 198)

[Establish Operational Levels from Historical Data](#) (see page 194)

Edit an Application Performance OLA

Edit a performance OLA to change the expected levels of service for a user-defined application on a group of networks that are identified by network type.

The management console creates a default OLA for each user-defined application to measure performance across all networks that do not have an assigned network type. In general, this is not recommended for monitoring Network Round Trip Time and Total Transaction Time because of differences in network latency between remote locations.


When you update the threshold values in a default OLA, your changes are dynamically applied to any new OLAs for that application that are configured to use the default OLA's value. For example, if the default OLA for the POP3 application requires at least 90% of the observations for Network Round Trip Time during a 1-hour interval to be at least 5ms, and you create an OLA for the POP3 application in the Austin network type, the default percentage of observations is 90% and the Network Round Trip Time threshold is 5 ms. If you change POP3's default OLA threshold for Network Round Trip Time, the management console dynamically updates the Network Round Trip Time threshold for POP3 Austin's OLA to use the new default value.

If you change the number of thresholds or minimum observations for an OLA metric without customizing the threshold values, the change applies to the default OLA.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Applications in the Show Me menu.
3. Scroll to the Application list, select the applications you want, and click OLAs.

Edit Application OLAs opens and the OLAs for the selected applications are listed by network type.

4. Click  to edit the network type that corresponds to the performance OLAs for a group of client networks.

(Optional) To edit the performance OLA for more than one remote location, select the appropriate network types, and click  to edit the selected OLAs.

Edit OLA Thresholds opens.

5. Edit the threshold settings for each OLA metric and click OK.

For information about setting server properties, click Help.

6. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

7. Wait a full reporting cycle, up to one hour, to see how the threshold changes compare to the actual response times observed by the management console.
 - a. On the Management page, in the Show Me menu, click Performance Detail OLA.
 - b. (Optional) Adjust the values for an operational level threshold. For example, if the threshold for Network Round Trip Time is less than 11 milliseconds at the 90th percentile, but only 67.744 percent of the observations meet this threshold level, adjust the threshold level percentile or the threshold level so that the application meets the operational level agreement for the remote site.

More information:


[Create an Application Performance OLA for a Group of Networks](#) (see page 196)

Delete an Application Performance OLA

Delete a performance OLA to remove any expected levels of service for one or more applications at a remote site. To delete an application performance OLA for the client networks at a remote location, identify the network type to which the client subnetworks are assigned.

Alternatively, you can remove a particular performance metric from an application performance OLA.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Applications in the Show Me menu.
3. Scroll to the Application List, select the applications you want, and click OLAs.
Edit Application OLAs opens and the OLAs for the selected applications are listed by network type.
4. Click  to delete an OLA.
Click Continue with Delete at the prompt to delete the specified OLAs.
5. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.
Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

More information:

[Group Client Networks by Network Type](#) (see page 49)

[Edit an Application Performance OLA](#) (see page 198)

Chapter 9: Managing Application Availability

This section contains the following topics:

[How Availability Monitoring Works](#) (see page 201)

[How Application Availability OLAs Work](#) (see page 207)

How Availability Monitoring Works

CA Application Delivery Analysis tracks the availability of user-defined applications through passive data observation. To determine application availability, the monitoring device observes IPv4-based client activity on each server that is assigned to an application.

Availability monitoring requires:

- CA Standard or Virtual Systems Monitor
- CA Multi-Port Monitor

If the monitoring device observes insufficient user traffic to an application port on a server during a 5-minute interval, the monitoring device verifies availability by making active requests to the application.

Availability criteria	What it means
Less than 10 successful TCP transactions	If there are less than 10 successful TCP transactions on the application port during a 5-minute interval, the monitoring device actively checks the availability of the application.
More than 10% refused sessions	If the application port refuses 10% (or more) connection requests during a 5-minute interval, the monitoring device actively checks the availability of the application.

The monitoring device actively checks application availability by sending a TCP SYN packet to the application port on the server. For an application defined by a port range, a connection is attempted on each of the first eight ports in the range.

If the monitoring device does not receive a SYN-ACK packet response from the application server:

- The application is rated Unavailable.
- The monitoring device actively checks the availability of the server that hosts the application by sending a ping request to the server. Server availability is not checked unless the application is rated Unavailable. If a server:
 - Responds to the ping, the server is rated as Available.
 - Does not respond to the ping, the server is rated as Unavailable.

For a load-balanced application, you can choose to rate application availability based on a minimum number of assigned servers rather than checking the availability of the application on each assigned server. For example, if the load balancer distributes application traffic between a minimum of 2 and up to 5 servers, it is normal for the application to be active on at least 2 servers. To rate the application as Available, acceptable application activity must occur on at least 2 of the 5 servers. If the load-balanced application is rated Unavailable, the [availability of each assigned server](#) (see page 206) is checked to determine which servers are available.

More information:

[Check Server Availability](#) (see page 206)

Why System-Defined Applications are Excluded

To avoid incorrectly opening an availability incident, the management console does not automatically monitor the availability of system-defined applications. For example, the first time the management console observes TCP traffic on Port 80 of a server, it will automatically create an instance of the HTTP application. But, enabling availability monitoring on that application would probably produce many false availability incidents, because the management console would expect to see HTTP running on every server on the domain, not just the server where the traffic was initially seen.

(Optional) To set an availability OLA for an application and apply the OLA to all of the servers monitored by the management console, create a user-defined application and assign it a domain. The management console automatically keeps the application server assignments up-to-date as your server subnets change.

How Server Incidents for Application Availability Work

The management console automatically opens a server incident in response to an application with an Unavailable status. If you have assigned an incident response to the:

- Server that hosts the application, and the application is Unavailable, the assigned incident response is launched, such as:
 - Email notification
 - SNMP trap notification
 - Ping Response Time Investigation
 - Performance via SNMP Investigation
 - Packet Capture Investigation
- Application, and the application is Unavailable, the assigned incident response is launched, such as:
 - Email notification
 - SNMP trap notification
 - Packet Capture Investigation

Note that the Application Connection Time investigation is not launched in response to an Unavailable application because the results of the investigation would not find any additional information about the status of the application.

How Application Availability Reporting Works

The management console reports the availability of the application when you click All Server Metrics in Settings. By default, Settings displays All Relative Metrics which does not display application availability status.

This rating is only applicable to a user-defined application where the management console administrator has assigned servers to the application. The management console reports:

- Whether the application was available on the server for the 5-minute interval.
- The percentage of time that an application was available.

For a load-balanced application, application availability is rated as the percentage of time that the application was available on a minimum number of servers.

More information:

[Check Server Availability](#) (see page 206)

Enable Availability Monitoring

The management console measures application availability across all client networks. You cannot monitor the availability of an application on a particular network.

Important! When using domains to separate tenant data, avoid false positive availability incidents by disabling availability monitoring. In an ISP environment, it is unlikely that a monitoring device can connect to an application server to actively check the availability of the application port.

Enable availability monitoring on the application, and optionally, on the servers that host the application.

Monitor Application Availability

To monitor application availability, edit the application properties to enable availability monitoring. By default, availability monitoring is disabled.

When deciding which applications to monitor for availability, we recommend monitoring:

- User-defined applications with assigned servers. The management console does not monitor application availability when the application is:
 - Automatically monitored (system-defined)
 - User-defined with an assigned server subnet
 - Assigned to all servers in a domain
- Priority Applications. If you set an availability OLA on an application that is not a priority application, and the management console grooms the application or filters the application data, the management console will unnecessarily check the availability of the application. The management console does not groom or filter priority applications.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Applications in the Show Me menu.
3. Scroll to the Application List, select the user-defined application you want, and click Edit.

Application Properties opens.

To more easily find user-defined applications, use the Show menu to hide system-defined applications.

4. Click the Availability Monitoring list and choose Enabled, then click OK.
5. If a load balancer distributes the application's traffic between servers, edit the application's server assignment to specify the number of servers that should be available.
6. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

More information:

[Check Server Availability](#) (see page 206)

[How Priority Applications Work](#) (see page 104)

Check Server Availability

The CA Application Delivery Analysis monitoring device checks server availability after it actively confirms that the application hosted by the server is unavailable. Checking server availability helps you determine whether an application availability issue is related to the application or the server which hosts the application.

If you do not want to actively check the availability of the servers that host the application, disable availability monitoring on each assigned server. By default, server availability monitoring is enabled.

To monitor the availability of a load-balanced application, in addition to checking server availability, you must also specify the minimum number of servers that must be available. For example, with a load-balanced application, the load balancer may not share the application load across all of its servers, so it is only necessary to verify that a minimum number of servers are hosting the application.

Note: Server availability is not checked when the application has a status of Available, therefore, it is possible for all of the assigned servers in a load-balanced application to have a status of Available while some of the servers are not actually available.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Applications in the Show Me menu.
3. Scroll to the Application List, select a user-defined application, and click Edit.
Application Properties opens.
4. Click Assignments.
5. (Optional) If the servers that host the application are part of a load-balanced farm, specify the minimum number of servers that must be available by following these steps:
 - a. Select Servers support the application in a load balanced farm. If this option is unavailable, click Properties and enable application availability monitoring.
 - b. Type a value for the Minimum servers required for the application to be considered available.
6. Click OK.
7. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

More information:

[Edit a Server](#) (see page 79)

How Application Availability OLAs Work

An availability operational level agreement (availability OLA) quantifies current application and server availability and availability trends in a report that can be presented to a broad audience. Knowing whether the server or application is not running is useful information. These scenarios point to different problems:

- An application is not running, but the server that hosts it is running.
- The TCP port is locked for this application only (such as TCP-80 for a Web application).

Availability OLA reports give management a tangible metric that shows the availability of an application and its assigned servers as opposed to only monitoring performance based on the number of outages or complaints received. The management console does not track network availability.

How Availability OLA Reporting Works

An availability OLA reports the percentage of time that an application is available across all networks. For example, an availability OLA can specify that an application that is hosted by a particular server is available to all networks 95.999% of the time over a one-month period. The management console updates availability OLA reports every 5 minutes to indicate whether an application meets the operational level agreement.

Enable an Application Availability OLA

Use the Availability OLA List to enable the availability OLA reporting for an application and specify a threshold for operational level compliance. When specifying a threshold, consider the time period you plan to use for reporting. For example, if you want to report on availability for a 12-month period, you may want to set a different threshold than you would use for reporting on a 1-month period. Note that you cannot create more than one availability OLA for the same application.

When you specify an availability OLA, make sure to also enable availability monitoring on the application itself. Specifying an availability OLA does not automatically update the application.

Follow these steps:

1. Click the Administration page.
2. Click Policies, Availability OLAs in the Show Me menu.

The Availability OLA List opens.

3. Click  to edit an OLA.

Configure Operational Level Agreement opens.

4. Select Enable Availability OLA, type the minimum percentage of time this application should be available, and click OK. You can specify a threshold value up to 3 decimal places, such as 93.999.

The management console displays a reminder to enable availability monitoring on the application.

More information:

[Edit a User-Defined Application](#) (see page 128)

Chapter 10: Managing User Account Permissions

This section contains the following topics:

[How User Account Permissions Work](#) (see page 210)

[Frequently Asked Questions](#) (see page 217)

How User Account Permissions Work

User account security is based on login access privileges and is fully compatible with that of the CA PC and CA NPC. The management console administrator creates and administers secure user accounts that are valid in the management console, the CA PC, and the CA NPC. These accounts allow other operators to access the management console administrative features and report data. Access to report data can be further restricted based on CA PC or CA NPC groups.

To create a secure system of access, authorized users gain access to product features based on the role and product privileges associated with their user account. When the management console is registered with the CA PC or the CA NPC, all tasks associated with creating and modifying these accounts are performed in the CA PC or CA NPC Administration pages.

The management console administrator might need to create additional user accounts to track the performance of selected components. For better security, the administrator should also plan to change the default password of the pre-configured administrator and user accounts.

The management console is required to be registered as a data source with either the CA PC or the CA NPC. The corresponding management tasks in the Administration pages of the management console are disabled. You manage users, roles, product privileges, and groups from the CA PC or the CA NPC.

Before setting up a management console user in the CA PC, familiarize yourself with the management console user roles, groups, and product privileges in the CA PC or the CA NPC. The CA PC or CA NPC security features—including users, roles, product privileges, groups, and domains—let you control which users can view specific data in the management console.

If you have not done so already, register the management console as a data source with the CA PC or the CA NPC. For information about registering the management console as a data source for the CA PC or CA NPC, and managing user security in the CA PC or CA NPC, see the online help.

When the management console is registered as a data source with the CA PC, the toolbar in the management console displays a CA PC hyperlink. When the management console is registered as a data source with the CA NPC, the toolbar in the management console displays a CA NPC hyperlink.

If you have registered the management console data source but you do not see the CA PC or CA NPC hyperlink, log out and then log back in to the management console.

Integrated Security

A *user account* specifies:

- The database credentials and authentication method that can be used to log into the management console. When registered as a data source with the CA PC or the CA NPC, the user account can be used to log into the management console and the CA PC or CA NPC.
- The management console role which defines the report pages of the management console, such as the Incidents page, that the user can access.
- The permission groups which define the report data, such as the data for a particular domain, that the user can access.
- The product privileges which define administrator-level permission, such as permission to access the Administration page, in the management console.

To manage user security on a management console data source, you must log into the management console with a user account that has the Administrator product privilege.

The default administrator account, `admin`, is locked to prevent changes to product privileges. This account is required to have Administrator privileges for all registered data sources. If you select a group of accounts that includes the `admin` account, you are not able to edit the product privileges for any of the selected accounts.

Use the following default user accounts or create your own. We recommend that you change the default passwords as soon as possible:

User Account	Default Properties	Default Password
<code>user</code>	Privilege: User Role: Network Operator Permission Group: All Groups Allowed to generate URLs: No	<code>user</code>
<code>admin</code>	Privilege: Administrator Role: Network Manager Permission Group: All Groups Allowed to generate URLs: Yes Use this account to manage management console users.	<code>admin</code>

User Account	Default Properties	Default Password
inv	Privilege: Power User Role: Network Engineer Permission Group: All Groups Allowed to generate URLs: Yes	inv

Product Privileges

A user must have product privileges on the CA Application Delivery Analysis data source to log in to the management console. Product privileges also specify access to the Administration page:

User

Gives access to all pages of the management console, except the Administration page.

Administrator

Gives access to all pages of the management console, including the Administration page.

Power User

Gives User-level product privilege, and Show Me menu access to the SNMP Profiles, Network Devices, and Device Groups on the Administration page.

Tip: If a user cannot log in to the management console user interface, verify that the user has been given a product privilege on the CA Application Delivery Analysis data source.

Roles

Role-based security enables a CA Application Delivery Analysis user to:

- Access portions of the management console.
- Drill into reports in the management console from a CA PC or CA NPC view.

More information:

[Drill into Data Sources Permission](#) (see page 213)

[Role Rights](#) (see page 213)

Role Rights

The following table summarizes role rights applicable to the CA Application Delivery Analysis (formerly NetQoS SuperAgent) management console:

Name of Role Right	Description
Engineering	Navigate the Engineering section; create Engineering reports
Operations	Navigate to the Operations section; create Operations reports
Management	Navigate the Management section; create Management reports
Incidents	Navigate the Incidents section; view Incidents reports
Investigations	Launch Investigations; drill into data from Investigations

Role rights do not give a CA Application Delivery Analysis user:

- Permission to access the Administration page of the CA Application Delivery Analysis management console.
To give a user access to the Administration page, give the user the Administrator or Power User product privilege on the CA Application Delivery Analysis data source.
- Access to actual report data in the CA Application Delivery Analysis management console.
To enable a user to see report data, assign the appropriate groups to the user.

More information:

[Product Privileges](#) (see page 212)

[Users and Groups](#) (see page 214)

Drill into Data Sources Permission

The Drill into Data Sources role right enables a user to drill down from a CA PC or CA NPC view into a supported data source, including CA Application Delivery Analysis. This role right applies to all data sources that are assigned to the user's role.

When configured with CA Multi-Port Monitor, this role right also enables a user to drill from CA Application Delivery Analysis into CA Multi-Port Monitor.

Important! The Multi-Port Monitor does not enforce permission sets from the CA PC or the CA NPC. For example, if a user has permission to a particular group of servers, and the Multi-Port Monitor can monitor at least one of the servers in the group, the Multi-Port Monitor displays performance data for all servers in the domain.

More information:

[Roles](#) (see page 212)

Users and Groups

The group-based security model in the CA PC and CA NPC lets you give access to the report data in the management console. By default, a CA PC or CA NPC role does not give a user permission to access the report data in a management console data source.

Create groups to give users access to some of the applications, servers, and networks, rather than all of the applications, servers, and networks in the management console data source. To enable a user to report on a group in the management console, you must give the user permission to access the group.

System groups let you give a user access to the management console data. If necessary, you can create custom groups to specify the data you want the user to access. For example, to enable a user to report on a system-defined application, create a server group that includes the servers that host the application port of interest.

System Groups

Create a system group to report on data from more than one management console or domain:

All Applications

Includes all applications reported by each management console data source.

All Domains

Includes group membership for each domain.

domain

Includes group membership for the domain. By default, the Default Domain is displayed.

All Servers

Includes all servers reported by each management console data source, and servers defined in other data sources.

Application Delivery Analysis Networks

Includes networks from all management console data sources.

Data Sources

Includes management console data sources that reported configuration information to the CA PC or CA NPC. A management console data source includes its own system groups.

If you assign a CA PC user account to the All Groups group, which includes all system groups, the user can see all the management console data in all reports that are included in her role selections.

Data Source System Groups

Data source system groups are automatically generated by CA PC and CA NPC and enable group-based reporting from both the CA Application Delivery Analysis management console and the CA PC or CA NPC management console on a particular CA Application Delivery Analysis data source. The following data source system groups are generated:

All Applications

Includes all applications reported by the data source, and organizes applications by type:

- Control Port Applications.
- FTP Applications.
- Standard Applications.
- Web Applications.

All Networks

Includes all client subnetworks reported by the data source, and organizes networks by:

- Domain. If you have not implemented domains to separate duplicate IP traffic, all are included in the Default Domain.
- Network Region Types.

All Servers

Includes all servers reported by the data source, and organizes servers by:

- Monitoring Device. Displays servers that are assigned to a monitoring device.
- Not Assigned. Displays servers that are not assigned to a monitoring device.

Relationships

Includes servers and applications, organized by relationship:

- Applications to Servers. Includes all servers assigned to an application.
- Servers to Applications. Includes all applications assigned to a server.

More information:

[Managing Tenants](#) (see page 95)

[Manage Servers](#) (see page 76)

[Group Client Networks by Network Type](#) (see page 49)

[Create a Control Port Application](#) (see page 125)

[Create an FTP Application](#) (see page 123)

[Create a Standard Application](#) (see page 117)

[Create a Web Application](#) (see page 122)

Tips for Groups

When you group applications, networks, and servers for reporting purposes, consider these issues:

- **Applications.** Only include multiple tiers of the same application in a single group. Trying to group application traffic that might be part of several applications at different tiers might yield unpredictable results. For example, the Telnet application might run on servers that are part of different multi-tiered applications throughout the Enterprise. Grouping all Telnet application traffic for a report would be misleading and make the management console data difficult to interpret.
- **Networks.** We recommend creating custom network groups. These groups are beneficial to compare similar networks or to see all networks with a particular bandwidth in reports. You could group networks of similar Network Round Trip Times because they often have similar bandwidths or you could group networks by functional departments such as Finance, Testing, and Development to compare operational levels.
- **Servers.** Server groups are useful when you troubleshoot specific applications. You could create an group including all servers used by the application to look for anomalous performance.

Frequently Asked Questions

- **Why can't a user log into the management console?**
If a user cannot log into the management console, make sure the user has [product privileges in the CA PC](#) (see page 212) or CA NPC to the management console data source.
- **After a user logs into the management console, why are some of the pages missing?**
If a user cannot view the Operations, Incidents, Management, or Engineering pages, make sure the user has the correct [role](#) (see page 212), and that the role has the appropriate selections. Access to each report page is controlled by the selections in the user's assigned role.
- **If the user cannot view the Administration page, make sure the user has [product privileges in the CA PC](#) (see page 212) or CA NPC to the management console data source.**
- **Why is there no report data in the management console?**
If a page in the management console does not display the appropriate data, make sure the user has [product privileges in the CA PC](#) (see page 212) or CA NPC to the management console data source.
- **Why can't I drill into the management console from CA PC or CA NPC?**
Your role must have the [Drill into Data Sources](#) (see page 213) role right, plus role rights to the appropriate pages of the management console data source.

Chapter 11: System Administration

This section contains the following topics:

[Windows Administrator Credentials](#) (see page 219)

[Manage the Database](#) (see page 219)

[Manage Console Settings](#) (see page 223)

[Change the IP Address](#) (see page 224)

[Manage SNMP Profiles](#) (see page 225)

[Manage Network Devices](#) (see page 229)

[Manage Scheduled Email](#) (see page 233)

[Perform System Maintenance](#) (see page 234)

Windows Administrator Credentials

CA appliances ship with the following default administrator accounts:

Windows 2008 operating system

netqos/Changepassword1

Windows 2003 operating system

netqos/changeme

nqadmin/qosisking

Manage the Database

The management console hosts a MySQL database for data storage and reporting.

Required Services

When operating normally, the management console automatically starts the services listed below.

Warning: To avoid data loss, do not attempt to manually stop or restart these services. For assistance, contact CA Support at <http://support.ca.com>.

- **CA ADA Availability Poller.** If the server that hosts an application is monitored by a CA Standard Monitor, the CA ADA Availability Poller service on the monitor checks the availability of the application. Otherwise, the CA ADA Availability Poller service on the CA Standard Monitor checks the availability of the application.
- **CA ADA Data Transfer Manager.** Synchronizes Cisco WAE devices to monitor application performance based on the applications, servers, and client networks defined on the management console.
- **CA ADA Inspector.** Loads 5-minute .dat files processed by the CA ADA Master Batch service into the database and communicates with the CA ADA Inspector Agent service to launch investigations.
- **CA ADA Inspector Agent.** If the server that hosts an application is monitored by a CA Standard Monitor, the CA ADA Inspector Agent service on the monitor launches investigations on the application, server, and related networks. Otherwise, the CA ADA Inspector Agent service on the CA Standard Monitor launches the investigations.
- **CA ADA Messenger.** Synchronizes monitoring CA Standard Monitor, CA Multi-Port Monitor, and CA GigaStor monitoring devices to monitor application performance based on the applications, servers, and client networks defined on the management console.
- **NetQoS MySql51 service.** Starts and stops the MySQL server which hosts the management console database.
- **CA ADA Monitor.** The CA ADA Monitor Service, located on the management console or the CA Standard Monitor, receives mirrored TCP packets and digest files from a CA GigaStor, Cisco WAE, or Cisco NAM device.
- **CA ADA Data Pump.** Performs weekly maintenance on the management console database.
- **CA ADA Master Batch service.** Receives data files from the CA ADA Batch service on the CA Standard Monitor for processing into 5-minute .dat files.

The Status of the Database

The Database Status page summarizes the amount of available disk space, the number of servers, applications, and networks that are currently monitored, and how fast the database is growing.

Follow these steps:

1. Click the Administration page.
2. Click Console, Database, Status in the Show Me menu.

Database Status opens.

3. View summary statistics about database growth.

For information about database summary statistics, click Help.

More information:

[How the Management Console Manages Database Growth](#) (see page 246)

Edit Database Storage Preferences

Edit the database storage preferences to specify:

- How long to keep report data in the database
- When to run the weekly database maintenance
- Who should receive email or SNMP trap notifications when available disk space falls below the specified threshold

A full or nearly full hard drive affects reporting of existing data and collecting new data. Several options are available for setting the duration of data storage and the type of data to retain. For more information, see the following sections. By default, the management console stores the following types of data for the following amounts of time:

Data Type	Stored for	Storage Time
Historical incident records	As long as 5-minute data is stored	1 to 12 months
5-minute data	1 month	1 to 12 months
Hourly data	6 months	1 to 12 months
OLA data	13 months	1 to 25 months

Follow these steps:

1. Click the Administration page
2. Click Console, Database, Maintenance in the Show Me menu.
Database Maintenance opens.
3. Edit the database storage preferences
4. Specify how the management console should notify you when available disk space falls below the threshold.

For information about setting database storage preferences, click Help.

5. Click OK.

Purge Data from the Database

The management console automatically maintains response time data along with your specified application, server, and network definitions, so it is typically not necessary to purge data. If necessary, you can purge all data and definitions, or purge a particular type of data, for example, 5-minute data, for a particular time period.

You cannot recover data after you purge. We recommend that you only purge data when requested to do so by CA Support.

Follow these steps:

1. Click the Administration page.
2. Click Console, Database, Purge Data in the Show Me menu.
The Purge Data dialog box opens.
3. Select an option to purge all data from the database or specify the time frame from which you want to purge a particular type of data, and click OK.

If the management console is registered as a data source with the CA PC or CA NPC , unregister the management console before you purge all data from the management console database. After the purge completes, re-register the management console as a data source with the CA PC or CA NPC. For information about unregistering the management console, see the CA PC or CA NPC online help.

4. Click Continue with Delete at the prompt to purge the specified data.

Back Up and Restore the Database

When installing the management console for the first time, the management console does not automatically perform a database backup as part of its weekly database maintenance.

Because various situations can lead to an unrecoverable database, we recommend that you schedule and perform weekly database backups. For more information, visit the CA Support website at <http://support.ca.com>.

Manage Console Settings

Use the management console settings to:

- View the name of the management console.
- Manage email settings.

When registered as a data source with the CA PC or CA NPC, the email settings are managed by the CA PC or CA NPC. If the management console is not registered as a data source, specify the following settings:

- IPv4 address of the SMTP server.
- Outbound TCP-25 access from the management console to the specified SMTP server.
- Reply To email address. Specify a valid return email address to prevent spam programs from filtering email notifications from the management console.
- Preferred chart image format for emailed reports: PNG or JPEG.
- Set the NRTT threshold to filter applications that use keep-alive messages.
- Verify the management IP address and order of the NICs, and send the management IP address of the management console to its monitoring devices.

Follow these steps:

1. Click the Administration page.
2. Click Console, Settings in the Show Me menu.
Console Settings opens.
3. Edit the console settings and click OK.
For information about console settings, click Help.

More information:

[Application Keep-Alive Messages](#) (see page 136)

[Change the IP Address](#) (see page 224)

Change the IP Address

Change the IP address of the management console to meet your needs. You must specify an IPv4 address in four-part, dotted-decimal notation.

When you change the IP address, the management console prompts you before automatically updating its monitoring devices to use the new IP address. Monitoring is temporarily interrupted while the management console-related services are restarted.


If the IP address of the monitor NIC is not routable, you do not need to change it.

Follow these steps:

1. Before you start, log out of the management console.
2. On the management console computer, log into Windows and complete the following tasks:
 - a. Update the management NIC properties to specify its IPv4 address in four-part, dotted-decimal notation.
 - b. In the Services Control Panel, restart the CA ADA-related services.
 - c. Log off.
3. In a Web browser, log into the management console.

The management console notifies you that the IP address of the management console has changed and prompts you to choose the IP address you want to use:

- a. Select the new IP address and choose the option to notify all monitoring devices of the new IP address.
- b. Click OK.
4. Verify the monitoring devices are communicating with the updated IP address of the Console:
 - a. Click the Administration page.
 - b. Click Data Monitoring, Monitoring Devices in the Show Me menu.
 - c. Scroll to the ADA Monitoring Device List and verify each monitor has a status of Running.

If necessary, click the blue gear menu  and click Set Console, then click Synchronize Monitor Devices to update the IP address of the management console on all monitoring devices.

Manage SNMP Profiles

A SNMP profile stores the SNMP user credentials required by the management console to:

- Query the SNMP agent:
 - On a server or network device for performance data, either in response to an incident on a server or network device, or as an investigation.
 - On a router to gather network definitions, or to gather performance information as part of a Trace Route investigation.

You can edit a SNMP profile to specify the port on which you want the management console to query the SNMP agent. By default, the management console queries the SNMP agent on UDP-161.

- Send a SNMP trap message in response to an incident on a server, application, network, or monitoring device.

Note that you cannot change the port on which the management console sends SNMP traps. The management console always sends SNMP traps to UDP-162.

The management console supports authentication to devices using SNMPv1, SNMPv2, and SNMPv3, and requires Read-Only permission.

When registered with the CA Performance Center or CA NetQoS Performance Center, the CA Performance Center or CA NetQoS Performance Center synchronizes SNMP profile changes between the management console and any other CA products that are registered with that instance of the CA Performance Center or CA NetQoS Performance Center. For more information about how the CA Performance Center and CA NetQoS Performance Center synchronize SNMP profiles, see the online help.

How SNMP Profile Discovery Works

When the management console performs a SNMP poll request, if the server or router is not assigned a valid SNMP profile, the management console discovers a valid SNMP profile by attempting to communicate with the SNMP agent using the list of available SNMP profiles.

When CA Application Delivery Analysis is registered as a data source with the CA PC or CA NPC, the CA Application Delivery Analysis Manager discovers a valid SNMP profile using the SNMP profiles synchronized from the CA PC or CA NPC that are also configured to be included in discovery.

When the management console discovers a valid SNMP profile, the management console assigns the SNMP profile to the server or network device. If the management console cannot discover a valid SNMP profile, the SNMP poll request times out.

The management console SNMP polls a server or network device:

- To perform a Performance via SNMP investigation
- As part of a Trace Route investigation
- To import network definitions from a router

Note that you can configure a SNMP profile to not be included in the discovery process. We recommend that you limit the number of SNMP profiles that can be used for discovery.

Add a SNMP Profile

Add a SNMP profile to enable the management console to query a server or network device, and to send a SNMP trap using a specified set of SNMP user credentials.

When the management console is registered as a data source with the CA PC or CA NPC, manage SNMP profiles in the CA PC or CA NPC. The CA PC and the CA NPC improves SNMP profile discovery by allowing you to specify the order where you want the management console, and any other registered data sources, to discover a valid profile.

We recommend that you limit the number of SNMP profiles that are included in the management console discovery process. For information about configuring SNMP profiles from the CA PC or CA NPC, see the online help.

Follow these steps:

1. Click the Administration page
2. Click Investigations, SNMP Profiles in the Show Me menu.

SNMP Profiles opens.

3. Click Add SNMP Profile.

SNMP Profile Properties opens.

4. Specify the SNMP profiles and click OK.

For information about setting SNMP profile properties, click Help.

More information:

[Managing Tenants](#) (see page 95)


[How SNMP Profile Discovery Works](#) (see page 226)

Edit a SNMP Profile

Use the SNMP Profiles list to view and manage the list of SNMP profiles between the management console and CA PC or CA NPC.

If you have registered the management console with CA PC or CA NPC, and the CA PC or CA NPC changes a SNMP profile, the changes are automatically synchronized with the management console. For information about how changes to a SNMP profile in the management console affect other data sources that are also registered with the CA PC or CA NPC, see the online help.

Follow these steps:

1. Click the Administration page.
2. Click Investigations, SNMP Profiles in the Show Me menu.
SNMP Profiles opens.
3. Click  to edit a SNMP profile.
SNMP Profile Properties opens.
4. Edit the SNMP profile properties and click OK.
For more information about setting SNMP profile properties, click Help.


Delete a SNMP Profile

Delete a SNMP profile to remove it from the management console. If you have registered the management console with CA PC or CA NPC, the CA PC or CA NPC synchronizes the change to any other registered data sources.

If a deleted SNMP profile was assigned to a server or network device, the management console unassigns the SNMP profile. When the management console needs to perform a SNMP poll request, and the server or router does not have an assigned SNMP profile, the management console discovers one.

If you delete a SNMP profile that is assigned to a SNMP trap notification, the management console automatically updates the SNMP trap notification to use a special SNMPv2 profile, *SuperAgent*, that is reserved for SNMP traps.

Follow these steps:

1. Click the Administration page.
2. Click Investigations, SNMP Profiles in the Show Me menu.
SNMP Profiles opens.
3. Click  to delete a SNMP profile.
4. Click Continue with Delete at the prompt to delete the SNMP profile.

More information:

[How SNMP Profile Discovery Works](#) (see page 226)

Manage Network Devices

We recommend adding a network device to the management console if you are planning to let the management console perform an SNMP query on that device, for example, to SNMP poll a router for its performance information.

If you do not assign an SNMP profile to a network device, the management console attempts to discover a valid SNMP profile.

The management console performs an SNMP query when:

- [SNMP polling](#) (see page 48) a router for network information.
- Launching a [Performance via SNMP investigation](#) (see page 172).
- Launching a [Trace Route Investigation](#) (see page 174).

More information:

[How SNMP Profile Discovery Works](#) (see page 226)

Add a Network Device

Add any type of network device to perform an investigation on that device.

Follow these steps:


1. Click the Administration page.
2. Click Investigations, Network Devices in the Show Me menu.
3. Click Add Network Device under the Show Me menu.
Device Properties opens.
4. Specify the network device properties and click OK or click Save and Add Another to add another device.

For information about setting network device properties, click Help.

View Network Device Investigations

Use the Network Devices List to view the results of any network device investigations.

Follow these steps:

1. Click the Administration page.
2. Click Investigations, Network Devices in the Show Me menu.
Network Devices opens.
3. Click  next to a device to open the Investigations Report page.
4. In the Investigations Report, specify the search criteria and click Search to filter the list of investigations.


The Investigations Report contents change according to your filter selections.

For information about setting search criteria, click Help.

Edit a Network Device

Edit a network device to update its properties, for example, by assigning an SNMP profile.

Follow these steps:


1. Click the Administration page.
2. Click Investigations, Network Devices in the Show Me menu.
3. Network Devices opens.
4. Click  to edit a device.
Device Properties opens.
5. Specify the network device properties and click OK or click Save and Add Another to add another device.

For information about setting network device properties, click Help.

Delete a Network Device

Delete a network device to prevent the management console from performing an investigation on that device.

Follow these steps:

1. Click the Administration page.
2. Click Investigations, Network Devices in the Show Me menu.
3. Network Devices opens.
4. Click  to delete a network device.
5. Click Continue with Delete at the prompt to delete the network device.

Group Network Devices for Investigation

Create a network device group to enable a management console administrator to schedule or launch an investigation on a group of network devices. For example, you can SNMP poll a group of routers for performance information.

Add a Network Device Group

To enable a management console administrator to schedule or launch an investigation on a group of network devices, add a network device group:

Follow these steps:


1. Click the Administration page.
2. Click Investigations, Device Groups in the Show Me menu.
Network Device Groups opens.
3. Click Add Device Group.
Device Group Properties opens.
4. Type a name in the Device Group Name box and click OK.
5. Add a device to the group by selecting devices in the Available Devices list and clicking the right arrow to move them to the Included Devices list.

Remove a device from the group by selecting a device in the Included Devices list and clicking the left arrow to move them to the Available Devices list.

View Investigations for a Network Device Group

Use the Network Device Groups list to view investigations that were run by a management console administrator on a group of network devices.

Follow these steps:

1. Click the Administration page.
2. Click Investigations, Device Groups in the Show Me menu.
Network Device Groups opens.
3. Click  next to the device group to view its investigations.
4. The Investigations Report opens.
5. (Optional) Specify criteria to filter the list of investigations:

Investigation Type

Type of investigation to view. Click All Investigations or the specific investigation type.

Target Type

Type of target for the investigation you want to view. Click All Targets or the specific target type (Device, Group, and so on).

Target


Target of the investigation you want to view. Select the specific target by name or IP address.

The Investigations Report contents change according to your filter selections.

Edit a Network Device Group

Edit a network device group to change the list of network devices that belong to the group.

Follow these steps:


1. Click the Administration page.
2. Click Investigations, Device Groups in the Show Me menu.
3. Network Device Groups opens.
4. Click  to edit a device.
Device Properties opens.
5. Specify the network device properties and click OK or click Save and Add Another to add another device.

For information about setting network device properties, click Help.

Delete a Device Group

Delete a network device group to prevent the management console from performing an investigation on the group of devices.

Follow these steps:

1. Click the Administration page.
2. Click Investigations, Device Groups in the Show Me menu.
Network Device Groups opens.
3. Click  to delete a device group.
4. Click Continue with Delete the prompt to delete the device group.

Manage Scheduled Email


A management console user's product privileges determine what a user can do with a report, for example:

- A user with permissions to view a report can create a schedule to send the report via email.
- A user with the Administrator product privileges can adjust the schedule or email settings for emailed reports, and schedule and send reports about monitoring devices.

Edit Schedules for Email Reports

View and manage the list of scheduled email reports, for example, to change email and scheduling options.

Follow these steps:


1. Click the Administration page.
2. Click Console, Scheduled Email in the Show Me menu.
3. Scheduled Email opens.
4. Click  to edit a scheduled email.
Scheduled E-Mail Properties opens.
5. Specify the scheduled e-mail properties and click OK.

For information about setting scheduled e-mail properties, click Help.

Delete a Scheduled Report

To cancel a scheduled report, delete the scheduled email.

Follow these steps:

1. Click the Administration page
2. Click Console, Scheduled Email in the Show Me menu.
3. Scheduled Email opens.
4. Click  to delete a scheduled report.
5. In the Delete Confirmation, click Continue with Delete to delete the scheduled email.

Perform System Maintenance

This section provides system administrators with information about how to set up and maintain CA Application Delivery Analysis.

How to Maintain Hard Disk Drives

The management console contains several tables that it consistently accesses for read/write operations. These tables occupy most of the disk space on the drive where the management console is installed. These I/O operations might cause disk fragmentation over time.

To maintain the hard disk drives, perform the following tasks:

- For databases larger than 10 gigabytes, defragment the D: drive or the drive on which the management console is installed each month.
 - Ensure that at least 20% of this drive contains free disk space before starting the defragmentation process.
 - Stop all CA ADA-related services, including the NetQoS MySQL51 service, before defragmentation.
 - Restart these processes after the defragmentation finishes.
- CA products are constantly writing to disk. During data seek for report compilations, the drive heads simultaneously write and read. As with any disk drive, this stress over time could lead to failure. To recover quickly with minimal data loss, schedule a regular database backup.

More information:

[Back Up and Restore the Database](#) (see page 223)

How to Update System Security and Install Windows Updates

The management console contains the latest Windows updates available at shipment time. Continue the maintenance of installing Windows updates and the latest anti-virus software. Because of varying security and administration policies, Microsoft's Automatic Update feature is disabled. Follow these steps to enable automatic update:

1. Enable Automatic Update through the Windows Control Panel.
2. Select the Notify me but do not automatically download or install any updates option. This option prevents automatic installation of an update that causes system instability.
3. Install important Microsoft updates. Do not install recommended or driver updates.
4. Reboot the server after you apply an update. If a third-party management system manages your updates, make the system automatically reboot after applying updates.

Warning: CA Application Delivery Analysis requires the Microsoft .NET Version 3.5 Framework. CA Application Delivery Analysis is not compatible with Microsoft .NET Version 4 or above. If a Windows update installs the .NET Version 4 or above Framework, CA Application Delivery Analysis may stop working. To correct this problem, uninstall .NET Version 4 or above and reinstall .NET Version 3.5 if necessary. Microsoft has issued the Knowledge Base Article [How to uninstall NET Framework 4](#).

Ensure Data Integrity and Use Anti-Virus Software

Follow the guidelines below to ensure data integrity:

- If you installed anti-virus software on the management console appliance, exclude all NetQoS directories from real-time and scheduled scans to avoid database corruption.
- If you configured your environment to push rules from a centralized anti-virus system, add a rule to exclude the following directories from scans:
 - C:\Windows\Temp
 - D:\NETQOS (or the directory where CA Application Delivery Analysis is installed) and all subdirectories.
- Automated backups can corrupt the database if the backup locks the database during a data write operation. If this occurs, manually restore the database.
- The management console does not support drive space compression. The benefit of greater drive space is not worth possible loss of the database or degrading system performance. See the Microsoft Knowledge Base Article [SQL Server databases not supported on compressed volumes](#) for more information. The compression-related issues discussed in this article are also applicable to MySQL databases.

More information:

[Back Up and Restore the Database](#) (see page 223)

Issues with Third-Party Software

Except for anti-virus, system management, and time-synchronization software, do not install third-party software, especially third-party network monitoring software such as *Wireshark*, on a Standalone management console or CA Standard Monitor. A third-party packet driver can interfere with packet monitoring and could void the warranty.

If you install third-party software on a Standalone management console or CA Standard Monitor, CA Support might ask you to uninstall this software before troubleshooting.

Issues with Domain Group Policies

The management console and the CA Standard Monitor appliances run the Windows operating system and can be added to a Windows domain. Depending on your environment, security policies and third-party software might be pushed to the CA Application Delivery Analysis appliance. Security policies and third-party software might cause problems with the normal operation of your appliance. Before you add a CA Application Delivery Analysis appliance to a Windows domain, exclude security policies, software, or both from being pushed to the server.

Product Upgrade Support

If you are running an earlier version of CA Application Delivery Analysis and have purchased a maintenance plan, download the latest version from the CA Support website at <http://ca.com/support>

Tip: Before you install a software upgrade or build, restart your server to ensure that the Windows updates take effect. Failure to restart can result in system failure and require an operating system rebuild.

Request a Hardware Replacement

If a CA Application Delivery Analysis appliance experiences a hardware failure, such as a hard disk, network interface card, or RAID controller, contact CA Support to request a replacement. CA Support verifies when you need to replace hardware and will require the following information:

- Server serial number
- Hardware component that failed
- CA Application Delivery Analysis software and version loaded on the server
- Operating system version loaded on the server
- Shipping address and personnel attention

When your replacement part or server arrives, a return shipping label is included. Repackage the failed part or server and place the return shipping label on the box.

Chapter 12: Monitoring Device Administration

This section contains the following topics:

[How Monitoring Devices Work](#) (see page 239)

[Create a Pair of Monitor Feeds](#) (see page 243)

[View Sessions Information](#) (see page 244)

[How the Management Console Manages Database Growth](#) (see page 246)

[Perform Basic Operations](#) (see page 252)

[Manage Monitoring Device Incidents](#) (see page 253)

[Troubleshoot Monitoring](#) (see page 261)

How Monitoring Devices Work

The management console lets you mix and match from any combination of embedded instrumentation from your Cisco switch infrastructure, dedicated monitoring and packet capture appliances, and WAN-optimization devices. This extensive monitoring architecture works without remote probes or agents, providing network groups a cost-effective approach to achieve their application delivery goals.

The management console consolidates response time data from the following types of monitoring devices:

Monitoring Device	Where to go for more information
CA Multi-Port Monitor	The <i>CA Multi-Port Monitor User Guide</i>
CA Standard Monitor	Monitoring with a CA Standard Monitor (see page 265)
CA Virtual Systems Monitor	Monitoring with a CA Virtual Systems Monitor (see page 301)
CA GigaStor	Monitoring with a CA GigaStor (see page 321)
Cisco WAAS	Monitoring with Cisco WAAS (see page 341)
Cisco NAM	Monitoring with a Cisco NAM (see page 371)
Riverbed Steelhead	Monitoring with Riverbed Steelhead (see page 387)

How Monitor Feeds Work

A monitoring device calculates response time metrics from one or more monitor feeds. A *monitor feed* is a source of response time data. For example, the:

- Packets monitor feed on a CA Standard Monitor receives mirrored TCP packets.
- Multi-Port monitor feed on a CA Multi-Port Monitor receives mirrored TCP packets.
- WAN Opt monitor feed on a CA Standard Monitor or CA Multi-Port Monitor receives packet digest files from Cisco WAE devices.
- Steelhead monitor feed on a CA Standard Monitor receives Steelhead-optimized packets.
- GigaStor monitor feed on a CA Standard Monitor or CA Multi-Port Monitor receives packet digest files from a CA GigaStor.
- NAM monitor feed on the CA Application Delivery Analysis Manager receives metric digest files from Cisco NAMs.

If more than one monitoring device observes traffic for the same server, the management console automatically assigns the monitor feed with the best source of response time data to a server.

If necessary, you can edit a monitor feed to assign:

- A secondary monitor feed. The management console automatically assigns the best monitor feed to a server. If you want the management console to automatically monitor a server from another monitor feed, for example, in the event that server migrates to a different location, assign a secondary monitor feed to a monitor feed.
- A domain. Assigning a domain to a monitor feed enables CA Application Delivery Analysis to uniquely identify server traffic. If you want to monitor duplicate IP traffic, assign a domain to each monitor feed.

More information:

[Create a Pair of Monitor Feeds](#) (see page 243)

[How Monitor Feed Assignment Works](#) (see page 241)

[Assign Domains to Monitor Feeds](#) (see page 100)

How Monitor Feed Assignment Works

The management console automatically evaluates server traffic from each monitor feed and assigns the most appropriate monitor feed to a server.

When a monitoring device initially observes a TCP session, the management console assigns the monitor feed that sees the highest inbound packet volume to the server. During the first hour, the assignment can change every 5 minutes depending upon the monitor feed with the highest packet volume. After an hour, the packet volume on a monitor feed must be significantly higher for the management console to automatically change the monitor feed assignment.

If the packet volumes are similar, the management console assigns the monitor feed with the faster Server Connection Time (SCT). Note that you can override the automatic monitor feed assignment and permanently assign a particular monitor feed to a server.

In a WAN-optimized environment, the management console assigns the best Packets monitor feed to monitor the Server segment and automatically calculates response time metrics from all WAN optimization devices.

If a server moves to a different location on the network, the automatically assigned monitor feed would automatically change to the monitor feed that now sees the server traffic. It can take up to 1 hour for the management console to change the monitor feed assignment.

To enable the management console to immediately observe traffic from more than one monitor feed, for example, as part of a load-balanced configuration or for disaster recovery and failover, you can assign a secondary monitor feed to a monitor feed.

More information:

[Create a Pair of Monitor Feeds](#) (see page 243)

[Edit a Server](#) (see page 79)

How Monitoring Device Synchronization Works

Synchronize monitoring devices to synchronize monitoring devices to monitor application performance based on the current client network, server subnet, and application definitions on the management console.

To minimize temporary interruptions to monitoring during synchronization, complete all changes before synchronizing monitoring devices.

If the monitoring devices are not synchronized, you are prompted to synchronize them:



Configuration Required

Close

Synchronize Monitor Devices. Settings on the Monitor Devices no longer match the Management Console.

Alternatively, you can [synchronize](#) (see page 252) a particular CA Standard Monitor or CA Multi-Port Monitor, or a CA GigaStor.

More information:

[Perform Basic Operations](#) (see page 252)

Create a Pair of Monitor Feeds

If you have load-balanced server traffic between 2 switches, or you have a primary switch and a secondary switch for failover purposes, create a pair of monitor feeds to enable the management console to automatically monitor assigned servers from both monitor feeds. With a pair of monitor feeds, the management console immediately reports the application traffic from either monitor feed.

Warning: Duplicate data will result if a pair of monitor feeds see the same traffic.


In the case of a failover, if you do not create a pair of monitor feeds and the server moves to a different location, it can take up to 1 hour for the management console to assign a new monitor feed and resume monitoring the applications hosted by the server.

For example, if you have two switches that are functioning as peers in a load-balanced configuration, the traffic between client networks and servers might be present on either switch, and the SPAN data from each switch is sent to a different monitoring device. If a server is associated with either one of these paired monitor feeds, the response time data for that server will be collected from both monitor feeds; in effect, there are two "best" monitor feeds for the server.


For WAN-optimized environments, the management console automatically computes response times from all WAN optimization devices, therefore, you do not need to assign a WAN Opt monitor feed to a server.

If you create a pair of monitor feeds, keep in mind that there will be differences in the observed response times between the primary and secondary monitor feed, for example, because of the location of the monitoring device and the server traffic that is available from that location, and changes in the proximity of the monitoring device to the server traffic.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
3. Scroll to the ADA Monitoring Device List and click  to edit a CA Standard Monitor or CA Multi-Port Monitor.

The Monitor Properties opens.

4. Scroll down to Monitor Feeds.
5. Click  to edit a monitor feed.
6. Click the Secondary Feeds list to assign a secondary monitor feed to:
 - The Packets monitor feed on a CA Standard Monitor

- The WAN Opt or GigaStor monitor feeds on a CA Standard Monitor or CA Multi-Port Monitor
- The Multi-Port monitor feed on the CA Multi-Port Monitor.

7. Click Update to apply your changes.

More information:

[Managing Tenants](#) (see page 95)

[How Monitoring Devices Work](#) (see page 239)

View Sessions Information

The management console displays information about the amount of IPv4-based TCP session activity that exists on a monitor feed, a monitoring device, or on all monitoring devices.

Use the active sessions information to verify that a monitor feed is monitoring TCP sessions. The management console reports the number of active TCP sessions that are observed by a monitor on a server, by application port. Child applications are treated as instances of the parent for this report; for example, two URLs for the same web application server would count as activity for the web server. It is not a count of the configured applications for the server.


View Active Sessions on a Monitor Feed

Use the Active Sessions page to view the number of active IPv4-based TCP sessions reported by a monitor feed during the last 5-minute reporting interval.

Alternatively, you can view active sessions information by monitor or management console from the last hour.

Note that the NAM monitor feed does not report active sessions information because the Cisco NAM Metric Agent computes its own response time metrics and does not include the TCP headers that are required to calculate sessions information.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
3. Scroll to the ADA Monitoring Device List and click  to edit the CA Standard Monitor or CA Multi-Port Monitor with the monitor feed you want.

Monitor Properties opens.


Use the Assigned To column to find the CA Standard Monitor or CA Multi-Port Monitor to which a Cisco WAE or CA GigaStor is assigned.

4. Click See Active Sessions under Monitor Feeds. Note that active sessions information is not applicable for the Cisco NAM monitor feed.

Active Sessions opens.

5. Click to expand a server and view its active sessions information.
6. (Optional) Click Email to email the active sessions information. For information about specifying email properties, click Help.

To send email, the [email settings](#) (see page 223) on the management console must be configured properly.

7. (Optional) Click the blue gear menu  and choose Reload Sessions to refresh the list with the active sessions information from the most recent 5-minute reporting interval.
8. (Optional) Click View Active Sessions for Feed and choose the monitor feed you want to view active sessions information for another monitor feed on the same monitoring device.

More information:

[View an Hourly Summary of Session Counts](#) (see page 246)

View an Hourly Summary of Session Counts


Use the hourly summary information to view the IPv4-based TCP session count by application, server, and monitor, for example, to:

- View the total number of sessions on each monitor, and view the session count on each server that is observed by a particular monitor.
- View the total number of sessions per application, and view the session count on each server that hosts the application.
- View the total number of sessions per server, and view the session count on each application that is hosted by the server.

The management console automatically updates the sessions information with observed IPv4-based TCP sessions from the last 5-minute interval. Note that sessions information is unavailable from a Cisco NAM monitoring device.

Reset the hourly summary information to begin reporting on the latest sessions information. Alternatively, you can view sessions information from the last 5-minute interval for a particular monitor feed.

Follow these steps:

1. Click the Administration page.
2. Click Console, Session Lists in the Show Me menu.
Sessions Lists opens.
3. Browse the lists of sessions by Monitoring Device, Application, and Server and expand a list to find the information you want.
4. (Optional) Reset the hourly summaries by clicking Reload Sessions.
5. (Optional) Send an email with the summary information in a list by clicking the blue gear menu .

More information;

[View Active Sessions on a Monitor Feed](#) (see page 245)

How the Management Console Manages Database Growth

Based on the server subnets, client networks, and port exclusions you specify, the management console monitors the busiest applications on each server. While the management console attempts to maximize its database resources, it is necessary to ensure that the database does not become unmanageable.

Database Capacity

The management console automatically limits the growth of its database to 120 million rows. A *row* contains the performance metrics for the TCP sessions between an application hosted on a server and a client network during a 5-minute interval. Note that the network portion of the combination can be the actual client IP address or, if you defined a network as a /23 mask (or lower), an aggregate of all the client IP addresses in the subnet.

Depending on the number of days you want the management console to keep 5-minute row data, the management console can monitor more or fewer combinations. By default, the management console keeps 5-minute data for 1 month which, on average, equates to approximately 4 million new rows per day to be stored in the database.

To better understand how the management console controls database growth, consider this example:

- A group of 254 accountants are all on the same /24 network, accessing the same application hosted on a single server.
- The accountants are at work 8 hours per day, but with lunch and breaks, an accountant is actually accessing the accounting application closer to 6 hours per day.
- In this case, the management console creates approximately 18,000 rows in the database per day.
[254 accountants x 1 application x 1 server x 6 hours x 12 rows per hour (once every 5 minutes) = 18,288 rows]

Following this example, the management console could monitor approximately 220 groups of accountants before it would become necessary to manage the size of the database.

Database Growth Control

To manage the 120 million row maximum, the management console:

- Filters each combination, using a minimum threshold for Network Round Trip Response Time (NRTT), to limit the impact of periodic client and server messages, such as keep-alives, on monitoring statistics in reports, and to control database growth.

If the number of NRTT observations during a 5-minute interval does not meet the minimum threshold, the management console does not create rows for that combination in the database.

From a reporting perspective, a filtered combination would prevent the management console from reporting session-level statistics for the filtered 5-minute interval, but the response time metrics for the application across the network are still valid.

If you do not want the management console to filter or groom an application, [prioritize](#) (see page 104) the application.

By default, the management console filters combinations for non-priority applications that have less than [10 NRTT observations](#) (see page 223) during a 5-minute interval.

- If the moving average for the row creation rate is too high, the management console limits the number of incoming rows by grooming low-volume applications that have the smallest number of application/server/network combinations. When the management console *grooms* an application, it does not create rows for the application/server/network combinations in the database. The management console does not groom priority applications or the busiest application on a server.

Grooming helps the management console maintain its database maximum of 120 million rows, and at same time, lets you report on the monthly time period that you want for the most interesting applications with the highest number of application/server/network combinations. The management console disables grooming when the moving average for row creation returns to an acceptable rate.

From a reporting perspective, the management console rates groomed application data as No Data (white).

If you do not want the management console to filter or groom an application, [prioritize](#) (see page 104) the application.

- If grooming does not successfully lower the moving average for the row creation rate, the management console does not create rows for non-priority application combinations with the smallest packet volumes. When the moving average for row creation returns to an acceptable rate, the management console resumes normal monitoring of non-priority applications.

From a reporting perspective, a filtered combination would prevent the management console from reporting session-level statistics for the filtered 5-minute interval, but the response time metrics for the application across the network are still valid.

If you do not want the management console to filter or groom an application, [prioritize](#) (see page 104) the application.

Monitoring Device Ratios

The monitoring device ratio varies by environment and depends on what you configure the management console to monitor. In particular, the number of client networks tends to have a significant impact on the amount of data produced.

If you are evaluating the management console, use the [daily database growth rate](#) (see page 221) from the evaluation to estimate the daily database row growth for production.

Keep in mind that on average, with default data retention settings, the management console can create up to 4 million rows per day. Under some conditions, a CA Multi-Port Monitor can generate 4 million rows per day, especially if the number of client networks is large. If the daily database row growth rate exceeds the threshold, the management console automatically manages the size of the database.

The processing load that is created by adding a monitoring device to the management console is expressed as a *monitoring unit*. For example, a CA Standard Monitor utilizes one monitoring unit. The management console supports up to 15 monitoring units.

The following list summarizes the equivalent processing load, in monitoring units, for each type of monitoring device:

Monitoring Device	Monitoring Units
CA Multi-Port Monitor	5
CA Standard Monitor	1
CA Virtual Systems Monitor	.5 (tier-to-tier traffic on the ESX host)
CA GigaStor	1 monitoring unit per 1 Gbps of SPAN
Cisco WAE (up to 50,000 optimized connections)	1
Cisco NAM-2 Blade (1 Gbps maximum)	1
Cisco NAM 2204 Appliance (2 Gbps maximum)	2
Cisco NAM 2220 Appliance (5 Gbps maximum)	5

Without any information about daily database growth, you can estimate that a management console can support one of the following monitoring device configurations:

- 30 CA Virtual Systems Monitor monitoring devices monitoring server-server traffic on an ESX host
- 3 CA Multi-Port Monitor appliances
- 10 Cisco NAM-2 blades and 1 CA Multi-Port Monitor
- 5 CA GigaStor appliances

Monitoring Device Recommendations

To optimize the management console database resources, follow the guidelines below:

- If possible, monitor from a physical switch using a physical monitor rather than from a virtual switch using a CA Virtual Systems Monitor. For example, monitor client-server traffic from the physical Distribution layer switch instead of the CA Virtual Systems Monitor on the ESX host that has the web server tier.

Physical switches have the best SPAN/VACL functionality and they do not suffer performance loss when you exercise that functionality. The CA Virtual Systems Monitor is a guest on the ESX host and as such, it consumes system resources. Also, the web tier is the most challenging tier to monitor since it deals with the client networks. Do not complicate ESX performance when you have a better choice from the physical realm.

- When monitoring from a physical switch using a physical monitor, monitor as close to the server as possible. When the monitoring device is close to the server, the improved accuracy of server metrics enables the management console to clearly distinguish between a network or server incident.
- If you have a CA Multi-Port Monitor, take advantage of SPAN aggregation on the monitor with hardware-based filtering to greatly improve efficiency.

Or, you can put a matrix switch between the mirrored switch ports and the monitor. This is an investment that can deliver big returns in a large environment. Use the matrix switch or network tool optimizer to filter unwanted traffic at line rate and send well-crafted packet streams to each port of a CA Multi-Port Monitor or CA Standard Monitor. This approach has several advantages:

- You can often adjust the data from the matrix switch rather than coordinate with the switch administrator to configure the physical switch during a change window.
 - Hardware-based filtering at line rate does not impact system CPU or memory.
 - Software-based filtering, which is available on the CA Standard Monitor or CA Virtual Systems Monitor, takes a lot of CPU and reduces monitor throughput.
 - Up-front cost and total cost of ownership are both lower.
 - The efficiency often makes it feasible to collect at the Access layer (closer to the servers) rather than just the Distribution layer.
- Sizing guidance is generic. Every environment is different. Rules of thumb are helpful, but actual capacity depends upon the nature of the application traffic as well as the configuration.

The management console does not place restrictions on the numbers of applications, servers, CPUs, or client networks monitored. The management console administrator is free to use the management console and monitoring devices to their maximum capacity, as dictated by the unique characteristics of each particular environment.

- Limit what the management console attempts to monitor by carefully defining:

- Application port exclusions.
- Server subnets.
- Client networks.
- Delete applications that are less interesting, like backup applications.
- Aggregate 24-bit client networks into broader subnets, like /22 networks. This approach limits the management console's ability to determine which TCP clients are affected by a performance degradation, but does reduce the number of rows that are created in the database.

More information:

[How Client Networks Work](#) (see page 29)


[Delete a System-Defined Application](#) (see page 115)

[Delete a User-Defined Application](#) (see page 129)

[How Port Exclusions Work](#) (see page 107)

[How Servers Work](#) (see page 67)

Perform Basic Operations

Perform basic operations on all of the monitoring devices in a list by using the blue gear menu (). Keep in mind that the:

- ADA Monitoring Device List includes the CA Standard Monitor, CA Virtual Systems Monitor, and the CA Multi-Port Monitor
- WAN Optimization Device List includes Cisco WAE devices
- CA GigaStor Device List includes CA GigaStor appliances

There are no basic operations to perform on a Cisco NAM.

Manage Monitoring Device Incidents

Manage monitoring device incidents to ensure that your monitoring devices are working properly. By default, the management console opens a monitoring device incident if:

- The management console does not receive data from a monitoring device for 15 minutes. For example, if a Cisco WAE is assigned to a CA Multi-Port Monitor, and the Cisco WAE stops generating response time data, after 15 minutes the management console creates a monitoring device incident for the Cisco WAE. Assuming the CA Multi-Port Monitor continues to process response time data on at least one of its logical ports, the management console does not create a monitoring device incident for the CA Multi-Port Monitor.
- A CA Standard Monitor or CA GigaStor cannot process more than 5 percent of the packets it receives.
- The management console cannot contact a CA Standard Monitor or CA Multi-Port Monitor for 5 minutes.

To notify the owner of a monitoring device about an incident, assign an incident response to a monitoring device.

When the condition that triggered a monitoring device incident no longer exists, the management console automatically closes the monitoring device incident. Therefore, you do not need to acknowledge monitoring device incidents.

Note that the CA Multi-Port Monitor performs some self-monitoring and can alert you to conditions that potentially affect its performance by sending SNMP trap notifications.

More information:

[Assign a Monitoring Device Incident Response](#) (see page 261)

View Monitoring Device Incidents

Use the Monitor Device Incidents list to view a summary of the monitoring device incidents, and display report details for a particular incident. Filter the list to choose the incident status and monitoring devices you want.

Follow these steps:

1. Click the Administration page.
2. Click Console, Monitoring Device Incidents in the Show Me menu.
3. Review the incident history in the Monitor Device Incidents list.

For information about incident status, click Help.

4. Filter the list of incidents you want by choosing an option:

Device

Choose a monitoring device from the list or use the default, All, to view all incidents for each CA Standard Monitor.

Incident State

Choose the incident states you want: Open, Closed, or Open and Closed.

Minimum Severity

Choose Major or Unavailable to filter the list of incidents. Unlike other thresholds, monitoring device incident thresholds have a severity of either Major or Unavailable.

5. Click a timestamp entry in the Incident on column to view the incident report details.

More information:

[Edit Monitoring Device Incident Thresholds](#) (see page 255)

[Enable and Disable Availability Monitoring](#) (see page 256)

Edit Monitoring Device Incident Thresholds

Set monitoring device incident thresholds to ensure that all of your monitoring devices are sending data to the management console. Depending on the type of monitoring device, you can configure additional thresholds for dropped or fragmented packets.

Note that you cannot adjust the incident threshold for monitoring device availability.

Specify monitoring device thresholds for:

- **Data inactivity.** If the management console stops receiving data from any type of monitoring device, the management console creates a Major (orange) monitoring device incident.

The management console considers a monitoring device to be *inactive* when it stops receiving performance data from the monitoring device. This can happen when:

- The network is down; no data is being generated.
- The monitoring device is down; data is generated, but the monitoring device is not active.
- The SPAN is lost; data is generated, but the SPAN is not active.

- **Packets dropped by the monitor.** If a CA Standard Monitor or CA GigaStor is too busy to process all of the packets it receives, and the packet capture driver drops too many packets, the management console creates a Major (orange) monitoring device incident. The management console does not monitor packet loss at the switch port or the Monitor NIC on the monitor. This threshold is only applicable to a CA Standard Monitor or CA GigaStor.


If the CA Standard Monitor or CA GigaStor consistently drops packets, [troubleshoot dropped packets](#) (see page 299).

- **Fragmented packets.** If a CA Standard Monitor or CA GigaStor receives fragmented packets, the management console creates a Major (orange) monitoring device incident. By default, this threshold is disabled. This threshold is only applicable to a CA Standard Monitor or CA GigaStor.

If fragmented packets occur consistently, one of the following conditions might exist:

- A malicious attack has occurred on the network.
- Devices on your network, such as routers or servers, might have improper maximum transmission unit (MTU) settings. Verify that you applied consistent MTU settings across the network to prevent fragmentation. If the MTU size is too large, retransmissions can occur if the packet encounters a router that is not able to handle that size. If the MTU size is too small, header overhead and the number of acknowledgments that need to be sent and handled increases.

Follow these steps:

1. Click the Administration page.
2. Click Policies, Performance Thresholds in the Show Me menu.
The Monitoring Device Thresholds List opens.
3. Click  to edit the thresholds for a monitoring device and click OK.

Monitoring Device Thresholds opens.


For information about setting incident thresholds for a monitoring device, click Help.

Enable and Disable Availability Monitoring

By default, the management console opens an Unavailable incident when the management console cannot contact a CA Standard Monitor, CA Virtual Systems Monitor, or CA Multi-Port Monitor for 5 minutes. To prevent the management console from opening Unavailable monitoring device incidents on a monitoring device, disable availability monitoring on the device.

Availability monitoring is not applicable for other types of monitoring devices. Note that you cannot specify a monitoring device threshold for availability.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
3. Scroll to the ADA Monitoring Device List and click  to edit a CA Standard Monitor or CA Multi-Port Monitor.

Monitor Properties opens.

4. Click Availability Monitoring to enable availability monitoring and click OK.

Add an Incident Response to a Monitoring Device

Add an incident response to enable the management console to launch a notification in response to a monitoring device incident. Unlike incident responses for applications, servers, and networks, a monitoring device incident response can be filtered by severity (Major or Unavailable) but not by duration.

By default, the management console does not launch a notification in response to a monitoring device incident.

Follow these steps:

1. Click the Administration page.
2. Click Policies, Incident Responses in the Show Me menu.
3. Click Add Monitoring Device Response under the Show Me menu.
Monitor Device Incident Response Properties opens.
4. Type an incident response name and click Apply.
5. Click Edit Actions in the third Show Me menu.
6. Click Add Action under the Show Me menu.
Monitoring Device Action Types opens.
7. Click an action and click Next.
Monitor Device Action Properties opens.
8. Complete the fields in Monitor Device Action Properties and click OK.
For information about setting action properties for a monitoring device, click Help.
The changes are automatically applied to the monitoring devices with the assigned incident response.

More information:

[Add an Action to a Monitoring Device Incident Response](#) (see page 259)

Edit Monitoring Device Incident Response Name


Edit a monitoring device incident response to rename it. Renaming a monitoring device incident response applies the name change to all monitoring devices with the incident response.

While you are editing a monitoring device incident response, you can also modify its responsive actions, for example, to edit or add an action.

Follow these steps:

1. Click the Administration page.
2. Click Policies, Incident Responses in the Show Me menu.

Monitoring Device Incident Responses opens.

3. Click  to edit a monitoring device incident response.

Verify at least one responsive action is assigned to each incident response, and that an incident response is assigned to each monitoring device.

4. Click Edit Incident Response in the third Show Me menu, .

Incident Response Name opens.

5. Type an incident response name and click OK.

More information:

[Add an Action to a Monitoring Device Incident Response](#) (see page 259)

[Delete a Responsive Action](#) (see page 260)


[Edit a Responsive Action](#) (see page 260)

Delete a Monitoring Device Incident Response

Deleting an incident response deletes the incident response and its responsive actions. If the incident response is assigned, the management console reassigns the default incident response to any affected applications, servers, or networks.

Before you delete an incident response, make sure the default incident response is set up properly or assign a new incident response to the affected applications, servers, or networks.


Follow these steps:

1. Click the Administration page.
2. Click Policies, Incident Responses in the Show Me menu.
The Monitoring Device List opens.
3. Click  to delete a monitoring device incident response. Note that you cannot delete the Default incident response for a network, server, or application.
The management console automatically reverts the monitoring device to the Default incident response.

Add an Action to a Monitoring Device Incident Response

Add an action to an incident response to enable the management console to launch one or more notifications or investigations in response to a monitoring device incident.

Follow these steps:



1. Click the Administration page.
2. Click Policies, Incident Responses in the Show Me menu.
Monitoring Device Incident Responses opens.
3. Click  to edit an incident response.
4. Click Edit Actions in the third Show Me menu.
5. Click Add Action under the Show Me menu.
Action Types opens.
6. Select an action and click Next.
Action Properties opens.
7. Complete the fields in Action Properties and click OK.
For information about setting action properties, click Help.
In Incident Response Actions, the new action is displayed.

Edit a Responsive Action

Edit a responsive action to modify the response to a monitoring device incident. By default, the management console does not launch a notification or investigation in response to a monitoring device incident.

Edit the default incident response to add one or more responsive actions, and create additional incident responses as needed.



Follow these steps:

1. Click the Administration page.
2. Click Policies, Incident Responses in the Show Me menu.
Incident Responses opens.
3. Click  to edit a monitoring device incident response.
4. Click Edit Actions in the third Show Me menu.
Incident Response Actions opens.
5. Click  to edit an action.
Action Properties opens.
6. Edit the responsive action settings and click OK. For more information, click Help.

Delete a Responsive Action

If you no longer want the management console to launch a particular responsive action, delete the action from its monitoring device incident response.

Follow these steps:

1. Click the Administration page.
2. Click Policies, Incident Responses in the Show Me menu.
Incident Responses opens.
3. Click  to edit a monitoring device incident response.
4. Click Edit Actions in the third Show Me menu.
Incident Response Actions opens.
5. Click  to delete an action.
6. In the Delete Confirmation, click Continue with Delete to delete the responsive action.

More information:

[Add an Incident Response to a Monitoring Device](#) (see page 257)


Assign a Monitoring Device Incident Response

In response to a monitoring device incident, the management console can send:

- An email notification
- A SNMP trap notification

Create an incident response with the responsive action you want, or assign an action to the default monitoring device incident response. The default incident response does not include an action.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
The monitoring device lists open.
3. Browse to a monitoring device list and click  to edit a monitoring device.
The Monitor Properties opens.
4. Click Incident Response to choose an incident response from the list and click OK.

Troubleshoot Monitoring

This section describes common problems you might encounter with monitoring TCP traffic and procedures for correcting these problems.

More information:

[How Monitoring Devices Work](#) (see page 239)

View the Monitoring Device Status

The CA Standard Monitor, CA Virtual Systems Monitor, and CA Multi-Port Monitor process incoming response time data for the management console. Use the ADA Monitoring Device List to view the status of a CA Standard Monitor, CA Virtual Systems Monitor, or CA Multi-Port Monitor, and to verify a monitor has recently processed incoming data from its monitor feeds.

If a monitor does not have a Running status, begin troubleshooting the monitor and any assigned monitoring devices such as a CA GigaStor.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.

The ADA Monitoring Device List opens.

3. Use the Status column to determine the monitor status:

Running

The monitor is receiving performance data from any of its monitor feeds.

Stopped

Indicates the management console is not currently receiving performance data from the monitor, but can contact the IP address of the monitor.

Unable to Contact


Indicates the management console cannot contact the management IP address of the monitor.

Never Contacted

Indicates that the management console has never contacted the monitor, and that you need to synchronize the monitor to collect performance data based on the client networks, server subnets, and applications that are currently defined on the management console. If necessary, use the Set Console command to update the monitor to communicate with the management console.

Stopped - No Data

Indicates that packet monitoring is disabled and the monitor is not receiving digest files from a Cisco WAE device, Cisco NAM device, or a CA GigaStor.

4. Use the Last Monitored column to verify the monitor has processed incoming data from its monitor feeds in the past 5 minutes.
5. To refresh the status information immediately, click the blue gear menu  and click Update Status.

More information:

[Monitoring Device Operations](#) (see page 279)

Troubleshoot Communication Issues

To report on application response times, the ADA Manager and its monitoring devices must be able to communicate with each other.

Verify	Can communicate with
ADA Manager	<ul style="list-style-type: none"> ■ The local MySQL database on TCP-3308 ■ Standard Monitor or Virtual Systems Monitor on TCP-1000, TCP-1001 and TCP-7878 ■ Multi-Port Monitor on TCP-80 and TCP-8080 ■ SSO on TCP-8381
Standard Monitor or Virtual Systems Monitor	<ul style="list-style-type: none"> ■ ADA Manager on TCP-80 and TCP-8080
Multi-Port Monitor	<ul style="list-style-type: none"> ■ ADA Manager on TCP-80 and TCP-8080
Cisco NAM Metric Agent on a Cisco NAM	<ul style="list-style-type: none"> ■ ADA Manager on TCP-9996
Cisco WAAS Flow Agent on a Cisco WAE device	<ul style="list-style-type: none"> ■ Standard Monitor on TCP-7878 ■ Multi-Port Monitor on TCP-7878 ■ ADA Manager on TCP-7878
GigaStor Connector on the GigaStor	<ul style="list-style-type: none"> ■ Standard Monitor on UDP-9995 ■ Multi-Port Monitor on UDP-9995 ■ ADA Manager on TCP-1001

Troubleshoot Missing Data

If the management console does not report on an application, server, or network, [verify](#) (see page 244) the Console sees TCP sessions for the application and server.

If the Console does not see the TCP sessions, verify:

- A port exclusion is not filtering out the application data. For more information, see [Application Port Exclusions](#) (see page 106).
- A [server subnet is defined](#) (see page 72) with at least one server that hosts the application traffic.
- A [client network is defined](#) (see page 37) with client IPs that communicate with the application.
- The monitoring devices are [synchronized](#) (see page 242) to monitor the applications, servers, and networks that are defined on the management console.
- If you have implemented domains, verify that you have assigned the correct domain to your monitoring devices, servers, and client networks.

Chapter 13: Monitoring with a CA Standard Monitor

This section contains the following topics:

[How a CA Standard Monitor Works as a Monitoring Device](#) (see page 265)

[Support for XFF Translation](#) (see page 270)

[Add a CA Standard Monitor](#) (see page 272)

[NAT Firewall Communication](#) (see page 274)

[Secure Packet Capture Investigation Files](#) (see page 275)

[Edit a CA Standard Monitor](#) (see page 276)

[Edit the Packets Monitor Feed](#) (see page 277)

[Manage Monitoring Device Performance](#) (see page 278)

[Monitoring Device Operations](#) (see page 279)

[Filter Out Keep-Alive Messages](#) (see page 281)

[Delete a CA Standard Monitor](#) (see page 283)

[Disable the Packets Monitor Feed](#) (see page 284)

[Troubleshoot a CA Standard Monitor](#) (see page 285)

How a CA Standard Monitor Works as a Monitoring Device

A CA Standard Monitor serves as a type of monitoring device for CA Application Delivery Analysis. The CA Standard Monitor passively monitors data center traffic from up to 2 ports and helps keep a continuous record of end-to-end system performance. A CA Standard Monitor resides on the same server as the management console; therefore, the management console is aware of the monitor. You can deploy additional monitoring devices to passively monitor mirrored TCP traffic from a server switch port. In addition, the CA Standard Monitor computes response time metrics for other types of monitoring devices, such as a CA GigaStor.

Note: For information about installing a CA Standard Monitor, see the *Installation Guide*.

How the CA Standard Monitor Works

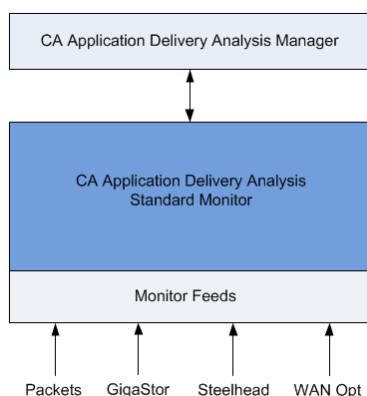
As shown in the following diagram, a CA Standard Monitor can receive and process performance data from several sources, including the:

- Packets monitor feed, which receives IPv4-based TCP packets on the monitor NIC from a SPAN or mirror port, or from a network tap, processes the data through the CA Standard Monitor, calculates and sends response time statistics to the management console for display in reports. We recommend local SPAN from the host switch for best results.
- GigaStor monitor feed, which receives packet digest files on the management NIC from a CA GigaStor, processes the IPv4-based TCP header information, and sends the response time statistics to the management console for display in reports.
- Steelhead monitor feed, which receives IPv4 Steelhead-optimized packets from a Steelhead appliance.
- WAN Opt monitor feed, which receives packet digest files from IPv4 Cisco WAE devices.

Based on the list of application ports, servers, and client networks you specify, the monitor creates response time statistics for all the matching server-application-network traffic it observes. The management console uses this information to automatically monitor the busiest TCP applications on each server. You can also define the applications you want to monitor and then load this configuration on the CA Standard Monitor.

The management console evaluates the response time metrics from all monitoring devices to assign the best monitor feed to each server, and to monitor the busiest TCP applications on each server.

If you want to configure a CA Standard Monitor to receive packet digests from a CA GigaStor, we recommend that you do not configure the monitor to simultaneously receive mirrored packets.



Note that a Standalone management console receives mirrored TCP packets on its Packets monitor feed.

More information:

[Managing Applications](#) (see page 103)

[How Monitor Feed Assignment Works](#) (see page 268)

Required Services

The CA Standard Monitor automatically starts the services listed below.

Warning: To avoid data loss, do not attempt to manually stop or restart these services. For assistance, contact CA Support at <http://support.ca.com>.

- CA ADA Monitor Management. Responds to requests from the management console to transfer .dat files from the monitor to the management console.
- CA ADA Data Transfer Manager. Synchronizes monitoring with Cisco WAE devices based on the applications, servers, and client networks defined on the CA Standard Monitor.
- CA ADA Inspector Agent. If the server that hosts an application is monitored by a CA Standard Monitor, the CA ADA Inspector Agent service on the monitor launches investigations on the application, server, and related networks. Otherwise, the CA ADA Inspector Agent service on the management console launches the investigations.
- CA ADA Messenger service. Synchronizes monitoring with the applications, servers, and client networks defined on the CA Standard Monitor.
- CA ADA Monitor. Located on the management console or CA Standard Monitor, this service receives mirrored TCP packets and digest files from its assigned monitoring devices, such as a CA GigaStor.
- CA ADA Batch. Stages .dat data files on the CA Standard Monitor for processing by the CA ADA Master Batch service on the management console.

How Monitor Feeds Work

A *monitor feed* is a source of response time data, for example, the:

- Packets monitor feed receives mirrored TCP packets on the monitor NIC.
- Riverbed WAN monitor feed receives WAN-optimized packets from a Riverbed Steelhead appliance.
- GigaStor monitor feed receives packet digest files on the management NIC from a CA GigaStor.
- WAN Opt monitor feed receives packet digest files on the management NIC from Cisco WAE devices.

The management console automatically assigns the monitor feed that is closest to the server as the source for monitoring TCP traffic on the server.

Edit a monitor feed to:

- Assign a particular domain. By default, a new monitor feed is assigned to the Default Domain. If you are not using domains to separate duplicate IP traffic, this is not applicable.
- Pair a secondary monitor feed for redundant data monitoring.
- View active sessions. Active sessions information helps you understand whether the monitor feed is monitoring active TCP sessions.

More information:

[Managing Tenants](#) (see page 95)

[Create a Pair of Monitor Feeds](#) (see page 243)

How Monitor Feed Assignment Works

If one of the monitor feeds on the CA Standard Monitor is the best source for monitoring TCP traffic on the server, the management console automatically assigns the monitor feed to the server.

More information:

[How Monitor Feed Assignment Works](#) (see page 241)

How Packet Capture Investigations Work

When the Packets monitor feed is assigned to a server, the management console performs packet capture investigations on the server from the corresponding CA Standard Monitor that sees the packets.

Unlike a CA GigaStor or CA Multi-Port Monitor, a CA Standard Monitor:

- Launches a packet capture investigation after the management console creates an incident.
- Runs one packet capture investigation at-a-time.
- Copies the packet capture file from the monitor to the user's local computer to view the packet capture investigation. Depending on the size of the packet capture file, it can take a long time to open the packet capture investigation.
- The CA Standard Monitor does not include support for long-term packet storage.

To enable a management console user to open a packet capture investigation created by a CA Standard Monitor, a network packet analyzer is required.

Monitoring Device Considerations

When using a CA Standard Monitor as a monitoring device, keep in mind:

- Typically, a CA Technical Representative helps you mirror TCP traffic to a monitoring device. For more information about mirroring TCP traffic, see the *Best Practices for Data Acquisition Guide*.
- PacketMon is the only packet sniffer certified for installation on a CA Standard Monitor or Standalone management console. To avoid conflicts with the packet capture driver, do **not** install any other packet sniffers, including *Wireshark*, on a CA Standard Monitor or Standalone management console.
- When configuring a packet capture investigation, keep in mind that a CA Standard Monitor does not include support for long term packet capture storage. To optimize available resources, configure packet captures to limit:
 - The maximum file size.
 - The number of bytes per packet.
- To limit access to the potentially sensitive contents of a packet capture, you can [disable packet captures](#) (see page 275) on the CA Standard Monitor.
- A CA Standard Monitor automatically monitors matching application traffic based on the client networks, server subnets, and port exclusions that are defined on the management console.
- A CA Standard Monitor monitors all types of applications.

Support for XFF Translation

When a user accesses a Web application through, for example, a proxy server, if the proxy server belongs to a different subnet than the clients it proxies, the management console incorrectly reports the Web application traffic from the proxy server's corresponding client network rather than the actual client network.

If the proxy server uses XFF, you can enable XFF translation in the CA Application Delivery Analysis Manager to extend monitoring support for Web applications and report the Web application traffic from the actual client network rather than the proxy server's corresponding client network.

Enabling XFF translation consumes additional resources on each CA Standard Monitor. By default, the CA Application Delivery Analysis Manager does not perform XFF translation.

More information:

[Create a Web Application](#) (see page 122)

[Enable XFF Translation](#) (see page 271)

How XFF Translation Works

The management console uses the X-Forwarded-For (XFF) HTTP header to identify a client's originating IP address when the client connects to a Web application through an HTTP proxy server that uses XFF. The typical XFF HTTP header format follows:

TCP Source IP: proxy3

X-Forwarded-For: client1, proxy1, proxy2

The list of IP addresses includes the farthest downstream client plus each successive proxy that passed the request and the proxy from which it received the request. In the example, the request passed proxy1, proxy2, and proxy3 (the closest proxy to the server).

Using XFF translation, the management console properly reports each client's Web traffic as belonging to the subnet that fits. In the management console, you will start seeing the proper volume of Web application traffic to and from the clients. You will also see a small amount of traffic on the proxy servers. The traffic consists of acknowledgements that the proxy server sends to the Web servers, not transactional traffic.

Tip: If you define the proxy server as its own network in a proxy environment, you can distinguish between performance on the Web server to proxy server link versus performance on the Web server to client subnet complete path. This is useful information if the proxy server is not co-located with the clients.

Enable XFF Translation

Enabling XFF translation requires you to run a MySQL command on the management console database. To run the required MySQL command, follow these steps:

1. Log into the management console server with a Windows administrator account.
2. Access the management console database using the correct database name and port:

Database name

The default management console database name is super.

Database port

The default port for the management console database is TCP-3308.

For more information about using MySQL commands, see the MySQL documentation available at www.mysql.com.

Follow these steps:

1. Log into the management console computer using a Windows administrator account.
2. Open a Command Prompt.
3. Login into MySQL by running the following command:
`mysql -P3308`
4. MySQL displays the following prompt:
`mysql>`
5. At the MySQL prompt, run the following command to switch to the management console database:
`use super;`
MySQL displays the following response:
Database changed
6. At the MySQL prompt, run the following command to enable XFF translation:
`INSERT IGNORE INTO parameter_descriptions (Parameter, Level, Type, DefaultValue, Description) VALUES ('XFFEnabled', 'System', 'boolean', '1', 'Non-zero to enable XFF endpoint extraction in URL monitoring.');`
MySQL displays the following response:
Query OK, 1 row affected (0.00 sec)
7. (Optional) To disable XFF translation after you enable it, run the following command:
`UPDATE parameter_descriptions SET DefaultValue='0' WHERE parameter='XFFEnabled';`
MySQL displays the following response:
Query OK, 1 row affected (0.02 sec)
Rows matched: 1 Changed: 1 Warnings: 0

8. Close the Command Prompt.
9. To apply your changes, [synchronize](#) (see page 279) the monitoring devices.

More information:

[Windows Administrator Credentials](#) (see page 219)

Add a CA Standard Monitor

Add a CA Standard Monitor to receive:

- TCP packets from a SPAN switch port.
- Packet digests from a [CA GigaStor](#) (see page 321).
- Packet digests from [Cisco WAE devices](#) (see page 341).

Prerequisites

Before you add a CA Standard Monitor, make sure:

- The [monitoring device ratio](#) (see page 249) is sized properly.
- The management console can communicate with the CA Standard Monitor on TCP-80, TCP-1000, and TCP-1001.
- Outbound UDP-161 is available, for example, to SNMP poll a server or network device.
- Outbound UDP-162 is available, for example, to send a SNMP trap.
- Outbound and inbound ICMP is available, for example, to verify that a server can respond to a ping request, and to measure the round-trip time.
- You can remotely access the CA Standard Monitor on TCP-3389 using Windows Terminal Services (RDP).

Add a CA Standard Monitor

When you add a CA Standard Monitor, the management console attempts to communicate with the monitor using the management IP address you specify. If you want to add a monitor that is not currently available on the network, the management console immediately polls the management IP address, after which you can specify the remaining monitor properties. When the monitor is available on the network, [synchronize](#) (see page 242) the monitoring device to establish communication between the monitor and the management console.

Optimize the available resources on the monitor by disabling TCP packet monitoring. For example, if you are planning to add a dedicated monitor to receive packet digest files from a CA GigaStor, and you do not want the monitor to monitor TCP traffic on its monitor NIC, disable packet monitoring.

After you add the monitor, it appears in the ADA Monitoring Device List. If you have defined [domains](#) (see page 95) in the CA PC or CA NPC, assign a domain to the Packets monitor feed on the CA Standard Monitor.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
3. Click Add ADA Monitor under the Show Me menu.
Standard Monitor Properties opens.
4. Complete the fields in Standard Monitor Properties and click OK.
For information about setting the monitoring device properties, click Help.
5. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

NAT Firewall Communication

To enable the CA Standard Monitor to communicate with a management console that is behind a NAT firewall, use the `LockConsoleAddress` utility to update the monitor configuration with the NAT address of its assigned Console.

This procedure is required, for example, when the management console is on a private network that is behind a NAT firewall. In this environment, it is possible for a CA Standard Monitor on the other side of the firewall to ping the management console on its NAT address, but the management console does not receive any reporting data from the monitor.

Before you run this utility, make sure you know the IPv4 NAT address of the management console to which you want to assign the monitor.

Follow these steps:

1. On the CA Standard Monitor, open a Command Prompt.
2. At the Command Prompt, change directories to the `<ADA_HOME>\bin` directory, for example `D:\NetQoS\bin`.
3. Type the following command and press Enter:
`LockConsoleAddress <NAT_IP>`

Where `NAT_IP` is the IPv4 NAT address of the management console. The utility updates the following registry key with the IPv4 NAT address you specify:
`HKEY_LOCAL_MACHINE\SOFTWARE\NetQoS\SACollector\Parameters\NAT_MasterDB`

4. (Optional) To verify the IPv4 NAT address that the monitor uses to communicate with the management console, type the following command and press Enter:
`LockConsoleAddress`

The results indicate the current IP address of the Console behind the NAT firewall.

5. (Optional) To delete the NAT address for the management console from the monitor, type the following command and press Enter:
`LockConsoleAddress -d`

The utility updates the following registry key to remove the IPv4 NAT address:
`HKEY_LOCAL_MACHINE\SOFTWARE\NetQoS\SACollector\Parameters\NAT_MasterDB`

Secure Packet Capture Investigation Files

A CA Standard Monitor stores its packet capture investigation files in unencrypted format. By default, a packet capture investigation only collects header information which greatly reduces the need for encryption.

To elevate security of packet capture investigation files:

- You can configure packet capture investigations to only capture header information.
- Disable packet capture investigations on the monitor. Note that when you upgrade the CA Standard Monitor, packet capture investigations are enabled. You must manually modify the monitor configuration after you upgrade to disable packet captures.

If you choose to enable packet capture investigations, you should also configure [roles](#) (see page 212) to limit who can create and view packet capture investigations.

Follow these steps:


1. On the CA Standard Monitor, open the Windows Explorer and browse to <ADA_HOME>\SuperAgent\dotnet\InspectorAgent.
2. In the InspectorAgent.exe.config file, uncomment the following entry:
<add key="Capture.CaptureTcp" value="disable" />
3. Restart the NetQoS Inspector Agent service on the monitor to apply your changes.
4. To manually delete existing packet capture investigation files on the monitor, navigate to <ADA_HOME>\SuperAgent\Web\batch\snifferfiles and delete the existing packet capture investigation (.enc) files. Note that when the 5-minute data is purged as part of database maintenance, the packet capture investigation files tied to that 5-minute data are automatically purged.

Edit a CA Standard Monitor

Edit the properties of a CA Standard Monitor to:

- Assign an incident response or to enable availability monitoring.
- View active sessions information for each monitor feed.
- View monitoring device incidents for the CA Standard Monitor.
- Edit the properties of any monitor feeds on the CA Standard Monitor.
- Perform basic operations on the CA Standard Monitor, such as synchronizing the monitoring device, stopping and starting, rebooting, and shutting down the monitor.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
3. Scroll to the ADA Monitoring Device List and click  to edit a CA Standard Monitor.

Standard Monitor Properties opens.

Note: The Standard Monitor type includes both the CA Standard Monitor and the CA Virtual Systems Monitor. If you are not sure whether the monitor is running on a virtual machine, verify its IP address.

4. Complete the fields in Standard Monitor Properties and click OK.

For information about setting monitoring device properties, click Help.

5. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.


Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

Edit the Packets Monitor Feed


The *Packets* monitor feed receives mirrored TCP packets from the monitor NIC on the CA Standard Monitor. Edit the Packets monitor feed to:

- Change the default name for the monitor feed.
- Assign a particular domain. By default, a new monitor feed is assigned to the Default Domain. If you are not using domains to separate duplicate IP traffic, this is not applicable.
- [Create a pair of monitor feeds](#) (see page 243) to enable the management console to automatically collect data for assigned servers from both monitor feeds.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
The ADA Monitoring Device List opens.
3. Click  to edit a CA Standard Monitor with an active Packets monitor feed.
Standard Monitor Properties opens.

Note: The Standard Monitor type includes both the CA Standard Monitor and the CA Virtual Systems Monitor. If necessary, verify the monitor IP address.

4. Click  under Monitor Feeds to edit the Packets monitor feed.
5. Edit the Packets monitor feed properties to assign a secondary feed or domain.
6. Click Update.
7. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

More information:

[Managing Tenants](#) (see page 95)

[How Monitoring Devices Work](#) (see page 239)

Manage Monitoring Device Performance


The management console automatically creates a Major monitoring device incident if a CA Standard Monitor:

- Stops sending data to the management console for more than 1 hour
- Does not process more than 5 percent of the packets it receives




More information:

[Edit Monitoring Device Incident Thresholds](#) (see page 255)

Monitoring Device Operations

Perform basic operations on some or all monitoring devices such as synchronize monitoring devices to collect performance data based on the client networks, server subnets, and applications that are currently defined on the management console. In the ADA Monitoring Devices List, use the blue gear menu  to perform basic operations.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
3. Scroll to the ADA Monitoring Device List and click the blue gear menu  to perform basic operations on all of the monitoring devices in the list.
4. (Optional) To perform a basic operation on a particular monitoring device, browse the list and click the edit icon . On the Standard Monitor Properties page, click the blue gear menu  and choose a command:

Start

Starts data monitoring on a monitor with a status of Stopped. While the monitor is Stopped, the monitor does not process response time metrics from any of its monitor feeds.

If the management console cannot communicate with the monitor, log onto the CA Standard Monitor to start the CA ADA Monitor service.

Stop

Stops data monitoring on a monitor with a status of Running. While the monitor is Stopped, the monitor does not process response time metrics from any of its monitor feeds.

Synchronize Monitor Devices

Updates the data monitoring configuration on a monitor to collect performance data based on the client networks, server subnets, and applications that are currently defined.

Reboot

Reboots a CA Standard Monitor or CA Virtual Systems Monitor. This command is not applicable to a CA Multi-Port Monitor.

If the management console cannot communicate with the monitor, log onto the monitor to reboot it.

Shut Down

Powers down a CA Standard Monitor or CA Virtual Systems Monitor. To power up a CA Standard Monitor or CA Virtual Systems Monitor after you shut it down, you must do so from the monitor. This command is not applicable to a CA Multi-Port Monitor.

If the management console cannot communicate with the monitor, log onto the CA Standard Monitor to shut it down.

More information:

[How Monitoring Device Synchronization Works](#) (see page 242)

[Perform Basic Operations](#) (see page 252)

Filter Out Keep-Alive Messages

The CA Standard Monitor offers an option to limit the impact of application keep-alive messages on monitoring statistics in reports. The technique involves limiting a selected application's Server Response Time (SRT) or Data Transfer Time (DTT) to a maximum value so that unnecessary SRT or DTT observations are ignored. You can set the value to a number of seconds that falls just below the keep-alive frequency.

If you suspect that an application is sending keep-alives, look for an inverse relationship between observations and SRT, and for SRT averages in the second range instead of the millisecond range. Configure the CA Standard Monitor to limit the maximum SRT for an application.

If you determine that your application uses keep-alives that result in extremely high Data Transfer Times (DTT), you can also apply a similar limit to filter DTT.

Use SRT and DTT filtering on the CA Standard Monitor with the Console's built-in [NRTT filtering](#) (see page 223). Typically, application keep-alives skew SRT data when clients are idle. NRTT filtering is useful at times when all clients are idle, such as non-business hours. However, during the day you can easily have a client network that has 10 idle connections and 30 active connections. You could surpass the NRTT threshold but still have skewed data due to the 10 idle connections inside that combination.

If you are unsure of the keep-alive frequency of a selected application, we recommend using 10 seconds as a safe starting point. It is highly unlikely that a server would take more than 10 seconds to start responding to a user request. In most (but not all) cases, the keep-alive frequency will be greater than 10 seconds.

For applications that use random ports, such as Microsoft Exchange 2007, the easiest way to identify the port is to open Outlook, run a netstat command on your computer, and record the dynamic port to which Outlook connects.

The required setting is also available on the CA Multi-Port Monitor. However, the steps to take to change the default setting are different. For more information, see the CA Multi-Port Monitor product documentation.

Follow these steps:

1. Access the CA Standard Monitor by using Microsoft Remote Desktop.
2. Browse to the installation directory, for example C:\CA\Bin.
3. Specify a SRT threshold for each port you want to filter:
 - a. Copy and rename the LimitServerResponseParams.ini.sav file to LimitServerResponseParams.ini.
 - b. In Notepad, edit the .ini file to specify the SRT threshold for each port you want to filter.

The LimitServerResponseParams.ini text file accepts multiple entries, separated by line breaks. Limit the SRT of each application by supplying the port number and the maximum amount of SRT that is allowed. Set the maximum SRT value to a value that is slightly less than the keep-alive frequency. For example, to ignore Citrix keep-alives that are occurring at a frequency of 60 seconds, supply the following entry:

```
/port=1494 /max seconds=59
```

- c. To filter a port range, uses the following syntax:

```
/min port=<lowerPort> /max port=<higherPort> /max seconds=59
```

If no port is specified, or is specified as being 0, then the specified limit is applied to all ports. A port range with 0 as its lower limit will have the same effect, regardless of what the upper port limit is specified as.


Entries that appear earlier in the file have priority over entries that appear later. Because of this, if there are various rules with overlapping port ranges, it is imperative to list the more specific ones first, otherwise they will be masked by the less specific ones. For example:

```
/port=23 /max seconds=15
```

```
/min port=100 /max port=200 /max seconds=50
```

```
/max seconds=120
```

This file limits SRT or DTT to 15 seconds for port 23, to 50 seconds for ports [100-200] and to 120 seconds for all other ports.

- d. Save the file.
4. If necessary, specify a DTT threshold for each port you want to filter:
 - a. Copy and rename the LimitDTTParams.ini.sav file to LimitDTTParams.ini.
 - b. Specify the DTT filter criteria as described in the previous step.
 - c. Save the file.
5. Open the management console and synchronize monitoring devices to apply your changes:
 - a. Click the Administration page.
 - b. Click Data Monitoring, Monitoring Devices in the Show Me menu.
 - c. Scroll to the ADA Monitoring Devices List and click the blue gear menu , then choose Synchronize Monitor Devices.

Delete a CA Standard Monitor


Delete a CA Standard monitoring device to remove it as a source of response time data. When you delete a monitoring device:

- Any servers that were pinned to the corresponding monitor feed are unpinned, and another monitor feed is automatically assigned. It can take up to 10 minutes to update the monitor feed assignment.
- Existing data is preserved for reporting purposes.

If you have pinned servers to a monitor feed, and you want to continue monitoring server traffic while the monitoring device is temporarily offline, consider the following options:

- Pin the servers to another monitor feed before you take the monitoring device offline. When you move the monitoring device back online, pin the appropriate servers to the monitor feed.
- Delete the monitoring device. Another monitor feed is automatically assigned, but it can take up to 10 minutes to update the monitor feed assignment.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
3. Click  to delete a CA Standard Monitor from the ADA Monitoring Devices List.

Note that the Standard Monitor type includes both the CA Standard Monitor and the CA Virtual Systems Monitor. If necessary, verify the IP address of the monitor.

4. In the Delete Monitoring Device Confirmation, click Continue with Delete to delete the monitoring device.

More information:

[Pinning a Monitor Feed to a Server](#) (see page 86)

Disable the Packets Monitor Feed

To optimize the available CA Standard Monitor resources for processing incoming digest files from Cisco NAM, Cisco WAE, Riverbed Steelhead, or CA GigaStor, disable packet monitoring on the CA Standard Monitor. Disabling packet monitoring removes the Packets monitor feed.

When you add a CA Standard Monitor, you can choose to disable packet monitoring. After you add a CA Standard Monitor, you cannot choose to disable packet monitoring. If necessary, delete the CA Standard Monitor and then add the CA Standard Monitor with packet monitoring disabled. If you have assigned any monitoring devices to the CA Standard Monitor, you will need to reassign them after you add the CA Standard Monitor.

More information:

[Add a CA Standard Monitor](#) (see page 272)

Troubleshoot a CA Standard Monitor

Troubleshoot the CA Standard Monitor to identify the cause of missing report data that should have been monitored by the CA Standard Monitor. If you are troubleshooting data monitoring issues with a Cisco WAE or CA GigaStor, make sure the CA Standard Monitor is receiving and processing the corresponding digest files.



Verify Active Sessions

Use the Active Sessions page to report on the number of active IPv4-based TCP sessions reported by each monitor feed on a CA Standard Monitor or CA Virtual Systems Monitor during the last 5-minute reporting interval. Keep in mind that a CA Virtual Systems Monitor only receives the Packets monitor feed.

If a monitor feed does not have any active sessions for a server or application, verify:

- The Packets monitor feed is seeing TCP traffic.
- The CA ADA Monitor service is running.

More information:

[View Active Sessions on a Monitor Feed](#) (see page 245)

[Troubleshoot the CA ADA Monitor Service](#) (see page 291)

[View Monitor Feed Statistics](#) (see page 287)

View Monitor Feed Statistics

On a CA Standard Monitor, the CA ADA Monitor service is responsible for calculating 5-minute averages from the incoming data it receives from each of its monitor feeds.

Use the monitor feed counters to help you better understand what the service is doing:

Packets Receiver

Packet data that is mirrored to the CA Standard Monitor.

GigaStor Receiver

Packet digest files received from a CA GigaStor. To display this counter, a CA GigaStor must be assigned to the CA Standard Monitor.

Steelhead Receiver

Packet digest files received from Riverbed Steelhead appliances.

WAN Opt Receiver

Packet digest files received from Cisco WAE devices. To display this counter, a Cisco WAE device must be assigned to the CA Standard Monitor.

Important: Before you begin, synchronize monitoring devices. The counter windows do not display until after you have synchronized data monitoring.

To view the monitor feed counters, you must log onto the CA Standard Monitor computer.

Follow these steps:

1. Log onto the CA Standard Monitor computer or use Microsoft Remote Desktop Connection (RDC) to remotely connect to the CA Standard Monitor computer.

When using Remote Desktop to connect to a Windows Server 2003-based server, use the /admin switch to connect to the physical console session. You must connect to the physical console session to view the feed receiver counters. For information about the /admin switch, see Microsoft KB 947723.

To view statistics for:

Packet data

Log onto the CA Standard Monitor or management console that receives the mirrored TCP packets.

Packet digest files from WAN optimization devices or a CA GigaStor

Log onto the CA Standard Monitor that receives the packet digest files.

2. Depending on the version of the operating system that hosts the CA Standard Monitor, the steps you need to take to view the monitor feed statistics vary. If the CA Standard Monitor is running on:

- Windows Server 2003, the feed receiver counters are automatically displayed. If they are not displayed, make sure you are connected to the physical console session.
 - Windows Server 2008, double-click the ADA Monitor Activity shortcut on the desktop to display the feed receiver counters.
3. If the counter descriptions are not displayed properly, close the counter windows and then reopen them by double-clicking the ADA Monitor Activity shortcut on the desktop.
 4. If the none of the feed receiver counters are displayed, verify:
 - The CA ADA Monitor service is running
 - The monitor is synchronized with the management console

More information:

[How Monitoring Device Synchronization Works](#) (see page 242)

[Monitoring Device Operations](#) (see page 279)

[View GigaStor Counter Statistics](#) (see page 339)

[View SPAN Receiver Statistics](#) (see page 289)

View SPAN Receiver Statistics

View Packets counter statistics for information about unoptimized TCP packet data that the CA Standard Monitor receives on a particular Monitor NIC.

Important: Before you begin, synchronize the monitoring device. The monitoring device must be synchronized to display its counter windows.

The Packets counter displays the following information:

span receiver 16 (Packets)	
Received Packets:	20,892,724,483
Dropped Packets:	27,574,614
To Server Packets:	6,895,724,466
From Server Packets:	7,162,452,301
To Server Bytes:	2,931,416,938,047
From Server Bytes:	5,281,006,374,641
Total Seen Packets:	14,066,496,745
Total Captured Bytes:	8,214,965,484,349
Accepted Sessions:	111,519,524
Rejected for Server:	0
Rejected for Client:	0
Rejected for Port:	0
Rejected for Positive:	0

Received Packets

Identifies the total number of packets that are received, but not inspected, by the Monitor NIC on the CA Standard Monitor.

Note: Inspection of the packet header allows you to use these packets to calculate application response time metrics. For more information, see Total Seen Packets.

Dropped Packets

Identifies the total number of packets that arrived at the Monitor NIC but the packet header was never inspected. If the CA Standard Monitor is too busy processing other packets and the packet capture driver buffer is full, the packet is dropped.

To Server Packets

Identifies the total number of packets that are sent from a client to a server.

From Server Packets

Identifies the total number of packets that are sent from a server to a client.

To Server Bytes

Identifies the total number of bytes sent from a client to a server.

From Server Bytes

Identifies the total number of bytes sent from a server to a client.

Total Seen Packets

Identifies the total number of packets that match a specified application port, client network, and server subnet.

Note: If the monitor is performing normally, the number of Total Seen Packets matches the number of Received Packets. If the monitor cannot inspect all the packets that it receives, the number of Total Seen Packets is less than the number of Received Packets, and the number of Dropped Packets increases. For more information, see Dropped Packets.

Total Captured Bytes

Identifies the total byte count for packets that match a specified application port, client network, and server subnet.

Note: The CA Standard Monitor inspects each packet header to determine whether the packet matches a specified application port, client network, and server subnet. For more information, see Total Seen Packets.

Accepted Sessions

Identifies the number of TCP sessions that match a valid application, server, and network combination on the management console.

Rejected for Server

Identifies the number TCP sessions where the server IP did not match a server subnet.

Rejected for Client

Identifies the number TCP sessions where the client IP did not match a client network.

Rejected for Port

Identifies the number TCP sessions where the server port matched the list of ports that the management console ignores.

Rejected for Positive

Reserved for future use.

Troubleshoot the CA ADA Monitor Service

The CA ADA Monitor service processes incoming monitor feed data. It:

- Loads, initializes, and controls the packet driver that reads each TCP packet traversing the monitor NIC on a CA Standard Monitor or management console.
- Processes the packet digest files from the management NIC, for example, from an assigned CA GigaStor.
- On the management console, it processes the metric digest files from a Cisco NAM.
- Reports on the TCP traffic that is observed by each monitor feed.

If the CA ADA Monitor service is Stopped, you can attempt to restart it. If you cannot start the CA ADA Monitor service, see the following sections for more troubleshooting information. This information applies to a CA Standard Monitor and to a management console that receives TCP packets on its monitor NIC or processes packet digest files from another type of monitoring device, such as a CA GigaStor.

Follow these steps:

1. On the Windows desktop, click the Start menu, and click Control Panel.
2. In the Control Panel, double-click Administrative Tools.
3. Double-click Services.
4. Right-click the CA ADA Monitor service and click Start.

Check the management console Log File

If the CA ADA Monitor service does not start, open the log file to find any error messages that can help you determine the root cause. By default, the log file is saved to <ADA_HOME>\Logs\SACollectorErrors[date].log.

Verify the Status and Priority of the NICs

If the CA ADA Monitor service is not starting, verify that only two NICs are enabled and that the management NIC has the highest priority.

Follow these steps:

1. On the Windows desktop, click the Start menu, and click Control Panel.
2. In the Control Panel, double-click Network Connections.
3. In the Network Connections window, ensure that the:
 - Management NIC and the Monitor NIC status is Enabled.
 - Status of all other NICs is Disabled.

Note: A status of Network Cable Unplugged also prevents the CA ADA Monitor service from starting.

4. To disable a NIC, right-click the NIC and click Disable.
5. In the Network Connections window, click Advanced Settings on the Advanced menu. Verify that the Management NIC appears first followed by the Monitor NIC and then other unused connections.
6. Start the CA ADA Monitor service. If the service does not start, follow the troubleshooting steps in the next section.

Verify the Registry Settings for Network Adapters

If the CA ADA Monitor service is not starting, verify the Windows Registry settings for network adapters.

Follow these steps:

1. On the Windows desktop, click the Start menu, then click Run.
2. In the Open box, type regedit.
3. Navigate to the following key:
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Class\{4D36E972-E325-11CE-BFC1-08002BE10318}

There should be several subkeys starting with 0000 and ranging to 0008. Each subkey represents a network adapter loaded in Windows.
4. Expand each subkey to ensure that it has a Linkage subfolder. If a subkey does not have a Linkage subfolder, export the subkey and delete it.
5. Export a subkey:
 - a. Select the subkey.
 - b. On the Registry Editor menu, click File and then Export.
 - c. Select a destination and file name.
 - d. Click Save.
6. Delete the subkey, right-click the subkey and click Delete.
7. Repeat steps 5 and 6 for all subkeys that do not have a Linkage subfolder.
8. Start the CA ADA Monitor service. If the service does not start, follow the troubleshooting steps in the next section.

Verify the NIC Settings

If the CA ADA Monitor service is not starting, use the `satstconsole.exe` program to verify the NIC settings.

Follow these steps:

1. In Windows Explorer, browse to `<ADA_HOME>\bin`.
2. Double-click `satstconsole.exe`.
3. In the Super Agent Console dialog box, note the Monitor IP address in the Adapter field. This is the IP address the management console is looking for.
4. In the Adapter field, type the IP address that is currently assigned to the monitor NIC.
5. Using the Microsoft Windows Network Connections window, verify that the Management NIC has a static IP address assigned to it.

Note: Problems occur if you use a DHCP address for the management NIC.

6. On the management console, click Start.
7. Start the CA ADA Monitor service. If the service does not start, follow the troubleshooting steps in the next section.

Verify the Registry Settings for the Monitor Service

If the CA ADA Monitor service is not starting, check the Windows Registry.

Follow these steps:

1. On the Windows desktop, click the Start menu, then click Run.
2. In the Open box, type regedit.
3. Navigate to the following key:
HKEY_LOCAL_MACHINE\SOFTWARE\NetQoS\SuperAgent\Parameters
4. Verify that the MasterDB key is set to the IP address of the management console or the CA Standard Monitor.
5. Verify that the Role key is set to *one* of the following:
 - CA Standard Monitor: Slave
 - Standalone management console: Standalone
6. Verify that the LastManagementAddress is set to the UNIX equivalent of the management NIC IP address of the CA Standard Monitor or management console.
7. To find the UNIX equivalent IP address, log into a server where MySQL is installed. Run the following query:
SELECT INET_ATON('x.x.x.x');
where x.x.x.x is the IP address of the Management NIC on the CA Standard Monitor or management console; for example:
mysql> SELECT INET_ATON('209.207.224.40');
8. Start the CA ADA Monitor service. If the service does not start, follow the troubleshooting steps in the next section.

Verify Communication with the management console

If the CA ADA Monitor service is not starting, verify that the CA Standard Monitor can communicate with the CA Application Delivery Analysis Manager on TCP-80.

The following procedure applies to a CA Standard Monitor that resides on the management console.

Follow these steps:

1. Open a command prompt on the CA Standard Monitor and type the following command:

```
telnet <host> <port>
```

where:

<host>

Is the IP address of the management console

<port>

Is 80

If you receive a "Connect failed" error, the port is blocked and needs to be opened. If you see a blank black screen, the connection was successful.

2. Start the CA ADA Monitor service. If the service does not start, follow the troubleshooting steps in the next section.

Verify At Least One Monitor Feed is Active

If the CA ADA Monitor service is not starting, verify that the Standard Monitor Properties are *not* configured to disable packet monitoring.

If the Disable Packet Monitoring option was selected when you added a CA Standard Monitor, you must assign a monitoring device, such as a CA GigaStor, to the monitor so that the CA ADA Monitor service will start.

Take a Packet Capture

Use PacketMon to take a packet capture of the TCP packets received by the monitor port on the CA Standard Monitor, and examine the packet header and contents.

PacketMon is the only packet sniffer certified for installation on a CA Standard Monitor or Standalone management console. To avoid conflicts with the packet capture driver, do **not** install any other packet sniffers, including *Wireshark*, on the monitor.

Follow these steps:

1. Install PacketMon on the monitor.
2. After the installation completes, start PacketMon.
3. In PacketMon, click Start to begin the packet capture.

If no results appear or if UDP/Unknown packet types appear, the SPAN is misconfigured or the network tap was not installed correctly. Verify the SPAN configuration or the network tap, and take another packet capture.

4. When valid TCP packets are returned, start the CA ADA Monitor service.

Check for Duplicate Client Networks

If you configured servers, applications, and client networks, make sure you have no duplicate client networks.

If you have a large configuration (several million combinations), the CA ADA Monitor service takes several minutes to start instead of starting instantly. If several hundred servers are defined, and you do not want to monitor all of them, remove them from the configuration. It is a best practice to reduce the management console configuration to only those servers, applications, and client regions that are of interest.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Networks in the Show Me menu.

The Network List opens.

3. Review the list of networks to make sure you did not configure the same client region twice.

Troubleshoot Missing Data

If the CA Standard Monitor is monitoring traffic, but the management console does not show data or the data is not current, there might be a problem in the communication between the management console and the monitor.

Follow these steps:

1. Verify that you can telnet from the management console to the CA Standard Monitor on TCP-8080.
2. If you cannot connect, verify that no firewalls or ACLs are preventing the communication. TCP-8080 is used by the management console to obtain data files from a CA Standard Monitor.
3. Verify that all CA ADA-related services are started on the management console and the CA Standard Monitor.
4. Verify that the <ADA_HOME>\Datafiles directory on the management console is not full of files that span several hours or days.

Troubleshoot Dropped Packets

A CA Standard Monitor will, occasionally, report dropped packets for a brief period of time, or extended periods of time. The length of time that this condition persists depends on several factors which can be identified by performing the following troubleshooting steps.

Follow these steps:

1. Evaluate the total amount of aggregate traffic in the SPAN to the monitor. The inbound data rate (and not the number of active sessions monitored by the monitor) could be causing the data to be dropped. This is a percentage utilization statistic for that port.

If the sheer amount of data traffic coming into the monitoring device is too high, reduce the amount of data that you SPAN to the Monitor port on the monitor. If the management console is configured to monitor all traffic, you may need to add a monitor and load-balance the server SPAN between the monitoring devices.

2. Evaluate the total number of active sessions in the Packets monitor feed. The monitor separates and calculates response time data for each active session, therefore, the more active sessions, the greater the impact to process new, incoming data.

The monitor creates an *active session* when it observes SPAN traffic that matches an application port, server, and client network that you have defined in the management console.

If you can update the management console to optimize the list of monitored application ports, servers, and client networks, you can optimize the processing resources on the monitor. If the management console is monitoring the right applications, servers, and network, then you may need to add a monitor and load-balance the server SPAN between the monitoring devices.

3. Evaluate background processes such as virus scans, spyware scanning and removal programs, and system backup processes. Each additional application running on the monitor impacts the total processor utilization and reduces the ability of the monitor to process inbound data traffic. The more memory, CPU, processor or disk IO bus-intensive these processes are, the fewer packets the monitor can actually process before dropping packets.

To optimize the available monitor resources, remove these applications from the monitor.

Chapter 14: Monitoring with a CA Virtual Systems Monitor

This section contains the following topics:

[How a CA Virtual Systems Monitor Works as a Monitoring Device](#) (see page 302)

[System Requirements](#) (see page 305)

[Add a CA Virtual Systems Monitor](#) (see page 306)


How a CA Virtual Systems Monitor Works as a Monitoring Device

The Response Time CA Virtual Systems Monitor (CA Virtual Systems Monitor) serves as a type of monitoring device for CA Application Delivery Analysis. Use the CA Virtual Systems Monitor to monitor IPv4-based TCP traffic between virtual servers on the same VMware ESX Host. The CA Virtual Systems Monitor monitors the server-to-server traffic between VMs on the same ESX Host, and minimizes the load on the CA Virtual Systems Monitor and by extension, maximizes the available resources on the ESX Host.

To monitor traffic from outside the virtual environment, for example, from another ESX Host or a remote subnet communicating with a virtual server inside the ESX Host, mirror the server traffic from the physical server switch to a physical monitoring device, such as a CA Multi-Port Monitor.

The CA Virtual Systems Monitor is only designed to monitor virtual server-to-virtual server communication "inside" the VMware ESX Host. Do not use the CA Virtual Systems Monitor to monitor traffic between physical (external) devices and a virtual server. To monitor physical-to-physical and physical-to-virtual communication, use a physical monitoring device, such as a CA Multi-Port Monitor.

As shown in the example below, the CA Virtual Systems Monitor monitors the server-to-server traffic between VMs on the same ESX Host, while physical monitoring devices take SPAN data from physical switches.



The CA Virtual Systems Monitor is similar to the CA Standard Monitor, but you install it on a VMware virtual machine rather than on a physical appliance. Like the CA Standard Monitor, the CA Virtual Systems Monitor passively gathers data from a mirrored port, inspects packet headers for performance-related information, and passes relevant performance metrics to the management console for reporting and display.

Unlike a CA Standard Monitor, a CA Virtual Systems Monitor does not receive digest files from Cisco WAE devices or a CA GigaStor.

The CA Virtual Systems Monitor supports SPAN (port mirroring) via the Cisco Nexus 1000V. If you are using the VMware vSwitch, the CA Virtual Systems Monitor requires a promiscuous port group to see mirrored traffic on the virtual switch.

More information:

[Monitoring with a CA Standard Monitor](#) (see page 265)

Plan for Deployment

The CA Virtual Systems Monitor must be able to "see" the virtual server-to-virtual server traffic on the ESX Host without also seeing any external traffic outside the ESX Host. Take the following into consideration as you plan for the installation:

- Which multi-tier applications do we need to monitor?
- Which servers on the ESX Host create server-to-server traffic? Enable the CA Virtual Systems Monitor to see the network traffic for this virtual server-to-virtual server communication.
- Which servers communicate with external devices, such as a virtual server on another ESX Host, or a physical server? Do not allow the CA Virtual Systems Monitor to see this communication. Instead, mirror the traffic to a physical monitoring device, such as a CA Multi-Port Monitor.

Note: Separating the external traffic at the monitoring device enables the management console to automatically assign the CA Virtual Systems Monitor to monitor the virtual server-to-virtual server traffic, and automatically assign a physical monitor to monitor the physical-to-virtual traffic.

- If the virtual machine you plan to monitor is assigned to more than one virtual NIC, the management IP address for the virtual machine must be assigned to a network adapter that belongs to the last physical NIC on the ESX Host (the virtual NIC). This is a VMware issue -- when selecting a network adapter, VMware only exposes the IP address of a network adapter if it is assigned to the last virtual NIC. For example, a virtual machine with 3 virtual NICs (VMNIC1, VMNIC2, and VMNIC3) will report the IP address for the network adapter that is assigned to VMNIC3, but not VMNIC1 and VMNIC2.

More information:

[Monitoring Device Recommendations](#) (see page 251)

Port Usage and Firewalls

When you set up the CA Virtual Systems Monitor and prepare to add it as a monitoring device, you need to consider any firewalls that could prevent communications between the CA Virtual Systems Monitor, which resides on a virtual machine, and the management console, which resides on a physical computer. You must determine:

- Which firewall ports are open?
- What types of traffic are allowed on those ports?

The CA Virtual Systems Monitor includes a Web service to communicate with the management console. The management console needs to send monitoring instructions to its monitoring devices periodically, and the CA Virtual Systems Monitor needs to send files that contain 5-minute aggregate data to the management console. Because the 5-minute data files consist of aggregated data, the network bandwidth that is consumed on the uplink port is minimal.

The following table summarizes the firewall ports that must be open to allow communications between the management console and the CA Virtual Systems Monitor:

Port	Direction	Description
TCP-1000	Inbound (from the management console to the CA Virtual Systems Monitor)	HTTP for management console access
TCP-80	Inbound	Web service requests for data
TCP-161	Inbound	SNMP MIB queries
UDP-162	Outbound	SNMP alert traps

System Requirements

The CA Virtual Systems Monitor must be installed on the Microsoft® Windows operating system. It is your responsibility to install a licensed version of the Windows operating system. CA does not include a license for the Windows operating system with the CA Virtual Systems Monitor.

To successfully deploy the CA Virtual Systems Monitor, your virtual machine should meet the following requirements:

Requirement	Description
Virtual switch	<ul style="list-style-type: none"> ■ VMware ESX® and VMware ESXi® 3.5 or 4.1 with the VMware vSwitch. ■ VMware ESX and ESXi 4.1 with the Cisco Nexus® 1000V.
Virtual machine to host CA Virtual Systems Monitor	<p>Processor: 1 virtual processor (<i>minimum</i>).</p> <p>Memory: 1 GB (<i>minimum</i>).</p> <p>Virtual Disk:</p> <ul style="list-style-type: none"> ■ Windows Server 2003: 16 GB ■ Windows Server 2008 R2: 30 GB <p>Network: 2 virtual network adapters (1 Gbit minimum). An uplink port (physical port leaving the ESX Host) is required to enable the CA Virtual Systems Monitor to communicate with its corresponding management console.</p>
Guest operating system	<p>The CA Virtual Systems Monitor requires either:</p> <ul style="list-style-type: none"> ■ Microsoft Windows Server 2003 Standard Edition (x64 or 32-bit) with SP2 ■ Microsoft Windows Server 2008 R2 Standard Edition (x64 only).

If you have registered the management console with CA PC or CA NPC, the CA PC or CA NPC reports virtual machine-related performance statistics from VMware along with performance data from the CA Virtual Systems Monitor.

To enable CA PC or CA NPC to display virtual machine-related performance statistics from VMware:

- VMware Tools must be installed on the VMs you plan to monitor.
- If the virtual machine you plan to monitor is assigned to more than one virtual NIC, the management IP address for the virtual machine must be assigned to a network adapter that belongs to the last physical NIC on the ESX Host (the virtual NIC). This is a VMware issue -- when selecting a network adapter, VMware only exposes the IP address of a network adapter if it is assigned to the last virtual NIC. For example, a virtual machine with 3 virtual NICs (VMNIC1, VMNIC2, and VMNIC3) will report the IP address for the network adapter that is assigned to VMNIC3, but not VMNIC1 and VMNIC2.

Add a CA Virtual Systems Monitor

To provision the CA Virtual Systems Monitor, complete the following tasks:

1. [Configure the virtual switch](#) (see page 306).
2. [Create the virtual machine](#) (see page 312).
3. [Configure the network connections](#) (see page 314).
4. [Run the CA Application Delivery Analysis setup program](#) (see page 317).
5. [Finish the setup](#) (see page 318).
6. [Post-Installation Steps](#) (see page 319).

Configure the Virtual Switch

The CA Virtual Systems Monitor works with the VMware vSwitch or the Cisco Nexus 1000V.

How to Configure the VMware vSwitch

To monitor the intra-VM traffic on the VMware vSwitch, create a dedicated Monitor port group with promiscuous mode enabled, and assign the Monitor NIC on the CA Virtual Systems Monitor to the Monitor port group.

The CA Virtual Systems Monitor is designed to monitor virtual server-to-virtual server communication "inside" the virtual environment. If you:

- Have separate vSwitches for front-end and back-end servers:
 - Enable promiscuous mode on the back-end vSwitch.
 - Mirror the front-end server traffic to a physical monitoring device.
- Do not have separate vSwitches for front-end and back-end servers:
 - Enable promiscuous mode on the vSwitch.
 - Mirror the front-end server traffic to a physical monitoring device and "pin" the front-end servers to the physical monitoring device.

To enable the Management network adapter on the CA Virtual Systems Monitor to communicate with the management console which is outside the ESX Host, configure the CA Virtual Systems Monitor to use an existing virtual network adapter or create a new Management port group. You do not need to enable promiscuous mode on the Management port group. For more information about promiscuous mode, see your VMware product documentation.


In the following example, the management console virtual machine (SA Collector) is configured to use the Monitor and Management networks. The ESX Host, 10.8.2.10, is configured with:

- vSwitch0. This virtual switch is assigned the vmnic0 device, and its security policy is configured to accept promiscuous mode on the virtual switch and the Monitor port group. Promiscuous mode enables the CA Virtual Systems Monitor to see the server-to-server traffic for the VMs that connect to the Monitor network, including the Manufacturing01 and SAP virtual machines.

The vmnic0 network adapter is the uplink port that enables client traffic from outside the virtual switch to communicate with a server on the virtual switch, such as an externally-facing Web server that sends client requests to the SAP application on the back-end SAP server.

To monitor front-end Web server traffic to external, physical clients (not shown), use a physical monitor, such as a CA Multi-Port Monitor, to monitor mirrored server traffic from the physical switch.

- vSwitch1. This virtual switch is assigned the vmnic1 network adapter, and the security policy for the virtual switch and the Management port group on the virtual switch are configured to reject promiscuous mode. The vmnic1 network adapter is the uplink port that enables the CA Virtual Systems Monitor to communicate with the management console, which is not hosted on the ESX Host.



Follow these steps:

1. On the VMware vSwitch, create a dedicated port group with Promiscuous Mode enabled:
 - a. Identify the VMware vSwitch that sees the application traffic you want to monitor.
 - b. On the VMware vSwitch, create a dedicated Monitor port group with the following settings:

Network Label

Name the network Monitor. Later, when you configure the virtual machine, you can easily identify the network adapter with the mirrored application traffic.

VLAN ID

Specify the VLAN ID with the application traffic you want to monitor or choose All. Leave this field blank if VLAN tags are not being used.

If you are on an ESX 3.5 Server and you want to monitor all VLANs, specify the VLAN ID as 4095.

If the virtual machine has an AMD adapter, you must configure the guest operating system to use the Intel E1000 driver. For more information, see <http://kb.vmware.com/kb/1004252>.

Promiscuous Mode

Configure the security policy to Accept promiscuous mode, which lets the Monitor port group to see all traffic on the VMware vSwitch.

2. To enable promiscuous mode on the Monitor port group, the VMware vSwitch and the Monitor port group must be configured to Accept promiscuous mode.

If promiscuous mode is not enabled on the VMware vSwitch configure the security policy on the VMware vSwitch to accept promiscuous mode.

3. If the vSwitch has an existing port group for management data, use this port group to enable the CA Virtual Systems Monitor to communicate with the management console.

If necessary, create a port group named "Management" to identify the network adapter that is not receiving the mirrored switch traffic.

How to Configure the Cisco Nexus 1000V

To enable the CA Virtual Systems Monitor to monitor the traffic on the Cisco Nexus 1000v, create a dedicated port profile for the CA Virtual Systems Monitor to see the appropriate VLAN traffic. After you provision the virtual machine, you can SPAN the VLAN traffic to the Monitor interface on the virtual machine. For information about mirroring VLAN traffic, see your Cisco product documentation.

When choosing the VLANs to monitor, if you:

- Have separate VLANs for front-end servers (traffic coming from outside the ESX) and back-end servers (internal traffic inside the ESX), mirror the front-end server traffic to a physical monitoring device.
- Do not have separate VLANs for front-end and back-end servers, mirror all the traffic in the ESX and mirror the external traffic coming to the ESX to a physical collector, then "pin" the front-end servers to the physical monitoring device.

Follow these steps:

1. In the Cisco Virtual Supervisor Module (VSM), create and enable a dedicated Monitor port profile that the CA Virtual Systems Monitor can use to see SPAN traffic for the appropriate VLANs. For information about creating a port profile, see your Cisco product documentation.

If there is an existing port group or port profile for management data, use this to enable the CA Virtual Systems Monitor to communicate with the management console. If necessary, create a port profile named Management so that when you configure the network connections on the CA Virtual Systems Monitor, you can easily identify the network adapter that lacks the mirrored switch traffic.

2. [Create the virtual machine](#) (see page 312) on which the CA Virtual Systems Monitor will be hosted, and configure the virtual machine to use the Monitor and Management network adapters.
3. SPAN the virtual server-to-virtual server VLAN traffic to the [Monitor interface](#) (see page 314) on the CA Virtual Systems Monitor virtual machine.

Additional SPAN Considerations

When configuring the SPAN, keep in mind:

- The source port can be an Ethernet port, a virtual Ethernet port or a VLAN interface. A standard SPAN session will only SPAN source and destination on the same physical (ESX) host.
- The destination port can be any physical or virtual Ethernet port. It cannot be a port channel.
- All destinations for a SPAN session must be on the same host. As an example, you cannot have one SPAN session with a VLAN as a source (on one host) and two destinations that are on different physical hosts.
- There is a limit of 16 SPAN sessions on the Cisco Nexus 1000V, but no more than 4 can have the same source (such as a VLAN). In the following example, both monitor sessions reference VLAN 152 as a source, therefore, only 2 more monitor sessions can be configured using VLAN 152 as the source.

```
monitor session 1
  source vlan 152 both
  destination interface Vethernet6
  no shut
monitor session 2
  source vlan 152 both
  destination interface Vethernet9
  no shut
```

Eliminate Duplicate Packets on VLANs

When the CA Virtual Systems Monitor receive two copies of a packet, for example, when you SPAN a VLAN to a CA Virtual Systems Monitor, the Packet Loss Percentage report on the Engineering page incorrectly reports an extremely high percentage of lost packets.

This section discusses how to enable the CA Virtual Systems Monitor to de-duplicate TCP packets.

Follow these steps:

1. Locate the RetransPacketDefs.ini.sav file in the C:\CA\bin directory.
2. Remove the .sav extension from the file name.
3. Edit the RetransPacketDefs.ini file. By default, the file contains the following entries:

```
<no logging>  
50 1000  
10 20 30 40 50 60
```

The first line tells the CA Virtual Systems Monitor where to log information about duplicate packets. If you replace the phrase <no logging> with a path to a logging file, such as C:\CA\bin\duppkts.txt, the CA Virtual Systems Monitor logs the information. Enabling the log file is a good practice and helps you determine whether the buffer size is adequate.

On the second line, the first number, 50, specifies that the CA Virtual Systems Monitor maintains a buffer of 50 packets to look for duplicates. If you reduce this parameter, the CA Virtual Systems Monitor consumes fewer CPU cycles looking for duplicates. This improves CA Virtual Systems Monitor performance, but possibly finds fewer duplicates. Note that the second number, 1000, is not used.

The last line in the file describes the bins of the histogram of duplicates, which appears in the log file specified on the first line. The histogram indicates how far each duplicate was found from its original. This information helps you tune the buffer size parameter. Ideally, lost duplicates occur in the first few bins and that the bin counts decline or disappear for the remote bins. If the counts for the high-numbered bins remain high, the buffer size is probably too small.

4. To apply your changes, restart the CA ADA Monitor service, wait 10 minutes for the management console to report de-duplicated data, and then check the Packet Loss Percentage report to verify that the packet loss percentage has decreased.

If necessary, increase the buffer size to enable the CA Virtual Systems Monitor to search for duplicates across a larger number of packets.

Create the Virtual Machine

Create a virtual machine to host the CA Virtual Systems Monitor. Use the CA Virtual Systems Monitor to monitor application performance between virtual servers on the same VMware ESX Host.

If you are using the Cisco 1000v for your virtual switching infrastructure, make sure you remember to SPAN the appropriate VLAN traffic to the Monitor NIC.

Make sure the virtual machine meets the following requirements.

Setting	Description
Virtual machine name	Specify a name for the virtual machine, for example, "CA Virtual Systems Monitor"
Data store	Select a data store with available space. You should plan to allocate 10 GB to the virtual machine.

Setting	Description
Guest operating system	<p>The CA Virtual Systems Monitor requires either:</p> <ul style="list-style-type: none"> ■ Microsoft Windows Server 2003, Standard Edition (64 or 32-bit) w/SP2 ■ Microsoft Windows Server 2008 R2 Standard (64-bit only). <p>After you install the operating system, do the following:</p> <ul style="list-style-type: none"> ■ Install Microsoft .NET Framework 3.5 SP1 on Windows Server 2003. Windows Server 2008 R2 already includes Microsoft .NET Framework 3.5 SP1. ■ Install SNMP and ensure that public is an accepted community name. ■ Install ASP.NET, which includes network COM+ access and IIS. When configuring IIS for installation, click to install the SMTP Service, and install the default IIS components (Common Files, Internet Information Services Manager, World Wide Web Service) and the IIS 6 Metabase Compatibility component. ■ Install any Critical or Important Updates. ■ Do not install Recommended Updates (also listed as <i>Software, Optional</i>). ■ Do not install Driver Updates (also listed as <i>Hardware, Optional</i>). ■ Recommended and Driver Updates should be cleared and the "Do not show me..." check box should be selected to hide the item from future scans. For assistance, contact CA Support at http://support.ca.com. ■ Uninstall Internet Explorer Enhanced Security Configuration on Windows Server 2003. ■ Install VMware Tools (Required).
Virtual CPU	Allocate 1 virtual CPU.
Memory	Allocate 1 GB (1024 MB)
NICs	<p>Create two network connections for the Management and Monitor NICs.</p> <ul style="list-style-type: none"> ■ Network Adapter 1, choose the Management network. ■ Network Adapter 2, choose the Monitor network.

Setting	Description
Virtual Disk	<ul style="list-style-type: none">■ On Windows Server 2003, create a new virtual disk and allocate 16 GB of virtual disk space. Be sure to choose the option to Allocate all space now. Under Advanced options, check Independent, and choose Persistent.■ On Windows Server 2008 R2, create a new virtual disk and allocate 30 GB of virtual disk space. Be sure to choose the option to Allocate all space now. Under Advanced options, check Independent, and choose Persistent.
Administrator accounts	<ul style="list-style-type: none">■ Specify a password for the administrator account, and select Password Never Expires.■ Create a new user in the local Administrator's group named <i>netqos</i> with a password of <i>changeme</i>, and select Password Never Expires.

More information:

[How to Configure the Cisco Nexus 1000V](#) (see page 309)

Configure the Network Connections

Configure the network connections on the guest operating system by completing the following steps:

1. [Rename the Network Connections](#) (see page 315).
2. [Configure the Advanced Settings for Network Connections](#) (see page 316).
3. [Assign IP Addresses to the Network Connections](#) (see page 316).

Rename the Network Connections

To more easily identify and configure the network connections you created when you configured the virtual switch, rename the network connections in the guest operating system to correspond to the Monitor and Management network adapters that are bound to the virtual machine.

Follow these steps:

1. In the VMware vSphere client, identify the MAC address of the Management and Monitor network adapters:
 - a. Right-click the virtual machine and click Edit Settings.
 - b. In the Virtual Machine Properties, make a note of the MAC address for the Monitor and Management network adapters.
2. In the guest operating system, identify the MAC address of each network connection:
 - a. On the Windows desktop, click the Start menu, right-click My Network Places, and click Properties. Or click Start, Control Panel. Right-click Network Connections and click Open.
 - b. In Windows Network Connections, right-click a connection and click Properties. In the properties dialog box, click the Support tab and click Details. The Physical Address corresponds to the MAC address of the Monitor and Management network adapters.

Tip: From a command prompt, you can run the `ipconfig /all` command to verify the MAC address.
3. Click Cancel.
4. In the Network Connections window, edit the default names to correspond to the appropriate interfaces as follows:
 - a. Management
 - b. Monitor

More information:

[Configure the Virtual Switch](#) (see page 306)

Configure the Advanced Settings for Network Connections

In the guest operating system, configure the advanced network connection settings to specify the correct connection order and bindings.

Follow these steps:

1. In Network Connections, on the Advanced menu, click Advanced Settings.
2. On the Adapters and Bindings tab, use the up arrow on the right side to move the Management NIC to the top. This action sets the priority and is necessary for CA Virtual Systems Monitor to operate properly.
3. In the Bindings box, clear the Internet Protocol (TCP/IP) option for all NICs on the following bindings:
 - File and Printer Sharing for Microsoft Networks
 - Client for Microsoft Networks
4. Click OK.

Assign IP Addresses to the Network Connections

In the guest operating system, assign an IPv4 address, subnet mask, and default gateway to the Management network connection.

If you are using the VMware vSwitch, configure the Monitor network connection with a non-routable IP address. When using a non-routable IP address, you do not need to specify a default gateway assignment.

Make a note of the IP addresses that you assign to the NICs. You will need this information to add the CA Virtual Systems Monitor to the management console.

Follow these steps:

1. On the Windows desktop, click the Start menu, and click Control Panel.
2. In Control Panel, click Network Connections.
3. Right-click Management, and click Properties.
4. On the General tab, click Internet Protocol (TCP/IP) and then Properties.
5. Click Use the following IP address, and type an IP address, subnet mask, and default gateway. Click OK.
6. Repeat these steps for the Monitor network connection, but specify a non-routable IP address and subnet mask, and do not specify a default gateway. We suggest the following values:

Monitor 1

IP address: 1.1.0.1

Subnet mask: 255.0.0.0

More information:

[How to Configure the VMware vSwitch](#) (see page 307)

Run the CA Application Delivery Analysis Setup Program

Run the CA Application Delivery Analysis setup program to install the CA Virtual Systems Monitor on the virtual machine. Before you run the setup program, make sure you can access the virtual machine by Remote Desktop.

The CA Application Delivery Analysis setup program is included with the CA Application Delivery Analysis .iso download file. If necessary, download the CA Application Delivery Analysis .iso from the CA Support website at <http://ca.com/support>.

Before you run the CA Application Delivery Analysis setup program, create a snapshot of the virtual machine so that, if necessary, you can uninstall the CA Virtual Systems Monitor. The CA Application Delivery Analysis setup program does not uninstall the CA Virtual Systems Monitor; you must revert to the previous (pre-installation) snapshot to remove it. For information about taking a snapshot of the virtual machine, see your VMware documentation.

Follow these steps:

1. Log into the virtual machine as a user with Administrator privileges.
2. Verify that the CA Application Delivery Analysis setup program has permission to run. To complete this step, right-click the setup executable file and click Properties. From this dialog box, click the Unblock button (if available) and click OK.
3. Run the CA Application Delivery Analysis Setup program. The setup program prompts you for the following:

End-user License Agreement

Accept the license agreement to proceed with the installation.

Installation path

Install the CA Virtual Systems Monitor to C:\CA.

Choose Type of Installation

Click Virtual Systems Monitor to install the CA Virtual Systems Monitor.

4. After the setup program completes, you must restart the guest operating system.
5. Log back into the guest operating system and [verify](#) (see page 289) that the CA Virtual Systems Monitor is seeing SPAN traffic.

Synchronize Time with a Time Server

If you have monitoring devices in different time zones, set each CA Virtual Systems Monitor to its local time zone and use a time server such as NTP to make sure the system time is accurate. Time is converted to Greenwich Mean Time (GMT).

Follow these steps:

1. Open a Command Prompt and run the following command:
`net time /queryntp`
Make a note of the SNTP server name.
2. Replace *<NTPServer>* in the following command with the SNTP server name returned in the query:
`net time /setsntp:<NTPServer>`
3. Configure Windows Time Service to start automatically.
4. Restart the computer.

Finish the Setup

Perform the following post-installation tasks:

1. Check the CA Support Web site at <http://support.ca.com> for software updates for the CA Virtual Systems Monitor.
2. (Optional) Configure the CA Virtual Systems Monitor to [filter duplicate packets](#) (see page 311).
3. Install anti-virus software, and exclude the following directories from scans:
 - C:\Windows\Temp
 - The installation directory, by default this is C:\CA, and all subdirectories.
4. Verify the system time and time zone. If you change these values, restart the computer.

Post-Installation Steps

After you complete the installation, you are ready to add the CA Virtual Systems Monitor to the management console. Keep the following points in mind:

- Add a CA Virtual Systems Monitor the same way you would [add a CA Standard Monitor](#) (see page 272). When adding the CA Virtual Systems Monitor, provide the [IP addresses](#) (see page 316) for the management and monitor NICs on the CA Virtual Systems Monitor.
- Managing a CA Virtual Systems Monitor is similar to managing a CA Standard Monitor. Use the ADA Monitoring Devices List to view and manage your CA Virtual Systems Monitor monitoring devices.

Note that in the ADA Monitoring Device List, a CA Virtual Systems Monitor has the same Type as a CA Standard Monitor--Standard Monitor. To distinguish a CA Virtual Systems Monitor from a CA Standard Monitor, you will need to know the host name or IP address of the management NIC on the CA Virtual Systems Monitor.

- Do not assign a Cisco WAE device or CA GigaStor to a CA Virtual Systems Monitor.

Chapter 15: Monitoring with a CA GigaStor

This section contains the following topics:

[How a CA GigaStor Works as a Monitoring Device](#) (see page 321)

[Add a CA GigaStor Monitoring Device](#) (see page 325)

[Block a CA GigaStor Input Port](#) (see page 332)

[Edit a CA GigaStor Monitoring Device](#) (see page 333)

[Edit the GigaStor Monitor Feed](#) (see page 334)

[Unassign a CA GigaStor](#) (see page 335)

[GigaStor Incidents](#) (see page 335)

[Perform Basic Operations](#) (see page 336)

[Delete a CA GigaStor Monitoring Device](#) (see page 336)

[Troubleshoot a CA GigaStor Monitoring Device](#) (see page 337)

How a CA GigaStor Works as a Monitoring Device

A CA GigaStor appliance (CA GigaStor) serves as a type of monitoring device for CA Application Delivery Analysis. CA Application Delivery Analysis identifies the time and source of a performance issue, and the CA GigaStor details root-cause analysis of the problem. A CA GigaStor:

- Is a constant-capturing device that captures all IPv4-based TCP packets. A single SPAN session with multiple destination interfaces enables a CA GigaStor and a CA Standard Monitor to share the SPAN. There is no need for a tap.
- Unlike a CA Standard Monitor, a CA GigaStor specializes in pure packet capture and can receive TCP packets on multiple interfaces (eight or more), from multiple SPAN ports.
- Captures the entire packet for complete conversation analysis and reconstruction.
- Combines TCP headers from multiple interfaces into a single GigaStor monitor feed.

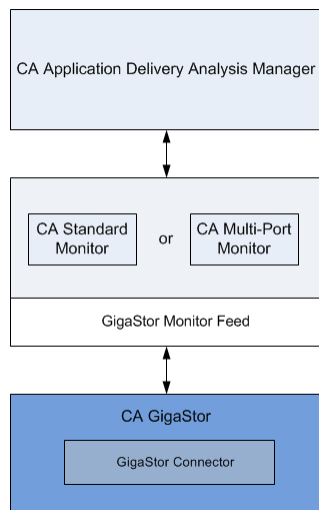
The CA GigaStor Connector enables a management console user to quickly and easily analyze large volumes of IPv4 packet data captured and stored on the CA GigaStor appliance, enabling network engineers to move seamlessly between the network, server, and application performance views in the management console and deep-packet inspection in CA Observer Expert for easier and faster troubleshooting.

How the CA GigaStor Connector Works

The CA GigaStor captures all packets that are passively copied to it, for example, from a SPAN or mirror port, or from a network tap, and captures the entire packet for complete conversation analysis and reconstruction.

The CA GigaStor Connector on the CA GigaStor polls the management console for a list of server subnets, client networks, and port exclusions. While the packets are in memory on the CA GigaStor, and before they are written to disk, the CA GigaStor Connector copies matching TCP header data into packet digest files and sends the digest files to its assigned monitoring device. If necessary, you can assign a CA GigaStor to more than one monitoring device. We do not recommend assigning a CA GigaStor to the CA Application Delivery Analysis Manager.

The CA GigaStor Connector sends packet digest files to its assigned CA Standard Monitor or CA Multi-Port Monitor for processing and calculation of response time metrics for each application, server, and network combination. The CA GigaStor Connector does not calculate response time metrics.



The packet digest files are received by the GigaStor monitor feed on the monitoring device. The management console then evaluates the response time metrics from all monitoring devices to assign the best monitor feed to each server, and to monitor the busiest TCP applications on each server.

You can also define the applications you want the management console to monitor and load this configuration on the CA GigaStor.

More information:

[Managing Applications](#) (see page 103)

How Monitor Feed Assignment Works

If a CA GigaStor is the best monitoring point for a given server, the management console automatically assigns the corresponding GigaStor monitor feed to that server.

If you have assigned a CA GigaStor to more than one CA Standard Monitor or CA Multi-Port Monitor for load-balancing purposes, the management console treats the GigaStor monitor feeds as a single "logical" feed for the purposes of server assignment. When manually assigning a GigaStor monitor feed, you can assign either GigaStor monitor feed to a server.

More information:

[How Monitor Feed Assignment Works](#) (see page 241)

How Packet Capture Investigations Work

When a CA GigaStor monitor feed is assigned to a server, the management console performs packet capture investigations on the server from the corresponding CA GigaStor that sees the packets.

When the management console launches a packet capture investigation on a CA GigaStor, the management console offers seamless drill-in to CA Observer Expert. CA Observer Expert offers packet-level details and expert analysis to identify the root cause of that performance issue.

You can launch or schedule a packet capture investigation on a CA GigaStor just as you would on another type of monitoring device. In the Packet Capture Investigation report, clicking the link in the Packet Capture Investigation report creates a CA Observer Expert filter to display the packets you want on the CA GigaStor in CA Observer Expert. Unlike a packet capture investigation on a CA Standard Monitor, there is no packet capture file to copy to your local computer.

Sizing Recommendations

A CA GigaStor sends packet digest files to the management NIC on its assigned monitoring device for processing. Depending on how the CA GigaStor is loaded, you should assign a CA GigaStor to at least one monitoring device. You cannot assign more than one CA GigaStor to a monitoring device.

To avoid overloading the Management NIC on the assigned monitoring device, avoid assigning a CA GigaStor to a monitoring device that is also receiving packet digests from another monitoring device, such as a Cisco WAE device.

Monitoring Device Considerations

When using a CA GigaStor as a monitoring device, keep in mind:

- Typically, a CA Technical Representative helps you mirror TCP traffic to a monitoring device.

Note: For information about data acquisition best practices, see the *Installation Guide*.

- When configuring the management console to perform a packet capture investigation on a CA GigaStor, you do not need to specify:
 - The maximum file size. The maximum file size is not applicable because the CA GigaStor captures and stores all packets, so there is no capture file. The packet capture investigation creates a CA Observer Expert filter to display the packets you want.
 - The number of bytes per packet. The CA GigaStor captures the entire packet.
- To limit access to the potentially sensitive contents of the capture, [grant a user permission](#) (see page 326) to a passive probe instance on the CA GigaStor.
- A CA GigaStor automatically monitors matching application traffic based on the client networks, server subnets, and port exclusions that are defined on the management console.
- The CA GigaStor Connector does not store packet digest files, therefore, if the CA GigaStor Connector cannot communicate with its assigned CA Standard Monitor or CA Multi-Port Monitor, the management console reports missing data. However, the CA GigaStor constantly captures TCP packets and you can use Observer to view the data on the CA GigaStor.
- The management console cannot monitor the performance of a Web application using the TCP-based performance data from the CA GigaStor Connector. The performance data from the CA GigaStor Connector does not include the HTTP header information required to determine the associated URL for a Web application.

To enable the management console to monitor a Web application with a CA GigaStor appliance, define a Standard application to monitor all TCP-80 traffic, or use CA Observer Expert to monitor the Web application.

Add a CA GigaStor Monitoring Device

To add a CA GigaStor as a monitoring device for CA Application Delivery Analysis, perform the following tasks:

1. [Verify the prerequisites](#) (see page 325), including required port access.
2. [Install and configure the required software](#) (see page 326) on the CA GigaStor appliance.
3. [Add the CA GigaStor monitoring device](#) (see page 327).
4. [Assign the CA GigaStor to a monitoring device](#) (see page 328).
5. [Install the required CA Observer software](#) (see page 326) on the user's client computer.
6. [Give the user permission](#) (see page 331) to access a passive probe instance on the CA GigaStor.

Prerequisites

Prerequisites for adding a CA GigaStor monitoring device are listed below:

- CA Observer software is installed on the CA GigaStor. If you are upgrading the Connector, you can download the CA Observer upgrade program from the CA Support website at <http://ca.com/support>.
- The time and date on the CA GigaStor are configured to match the time and date on its assigned CA Standard Monitor or CA Multi-Port Monitor.

The performance data collected by a CA GigaStor is time stamped using the timestamp of its assigned monitor. If you configured the monitor to use the Network Time Protocol (NTP), configure NTP on the CA GigaStor.
- The [monitoring device ratio](#) (see page 249) is sized properly.
- A CA Standard Monitor or CA Multi-Port Monitor is available to process digest files from the CA GigaStor Connector into response time data. To maximize the available resources on a CA Standard Monitor, disable packet monitoring when you [add the CA Standard Monitor](#) (see page 272).
- The management console can communicate with the CA GigaStor on TCP-1001.
- The CA GigaStor can communicate with its assigned CA Standard Monitor or CA Multi-Port Monitor on UDP-9995.
- The management console users who want to open a packet capture investigation on the CA GigaStor can communicate with the CA GigaStor on TCP-25901 through TCP-25903.
- You can access the CA GigaStor by Remote Desktop with Windows Terminal Services (RDP) on TCP-3389.

Install and Configure Software on the GigaStor Appliance

Install the required software by performing the following tasks:

- Install the CA GigaStor Connector on the CA GigaStor appliance.
- Create a passive probe instance on the CA GigaStor for each management console user.

Install the CA Gigastor Connector

Install the CA GigaStor Connector on the GigaStor appliance to use your CA GigaStor appliance as a monitoring device for CA Application Delivery Analysis. The version number of the Connector should match the version of CA Observer. Contact CA Support to download the correct version.

To install the GigaStor Connector, run the Connector setup on the GigaStor desktop. Reboot when setup completes.

Create a Passive Probe Instance for Each User

To enable a management console user to open a packet capture investigation on a CA GigaStor, create a passive probe instance for each management console user you want to access the CA GigaStor.

To avoid accidentally disconnecting a management console user from CA Observer Expert, do not share passive probe instances between management console users. If a management console user opens a CA GigaStor packet capture investigation using a CA GigaStor probe instance that is already in use, the existing user is disconnected from the CA GigaStor.

For information about creating passive probe instances on the CA GigaStor, see the CA GigaStor product documentation.

Add a CA GigaStor Monitoring Device

Add a CA GigaStor monitoring device as a data source for the management console. If you need to separate duplicate IP traffic by domain, after you add the CA GigaStor, [edit the GigaStor monitor feed](#) (see page 334) and assign the GigaStor monitor feed to the domain you want.

After you add the CA GigaStor, [assign the CA GigaStor to a monitoring device](#) (see page 328).

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
3. Click Add GigaStor under the Show Me menu.
GigaStor Properties opens.
4. Complete the fields in GigaStor Properties and click OK.
5. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

Assign a CA GigaStor to a Monitoring Device

After you add a CA GigaStor to the management console, assign it to a CA Standard Monitor or CA Multi-Port Monitor to process digest files from the CA GigaStor Connector into response time data. Do not assign a CA GigaStor to a CA Virtual Systems Monitor.

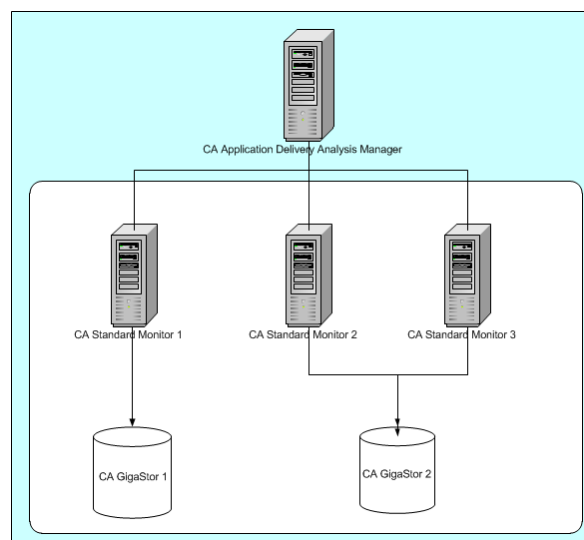
To maximize the available resources on the CA Standard Monitor, disable packet monitoring when you [add a CA Standard Monitor](#) (see page 272).

To load-balance the monitor feed from a single CA GigaStor for scalability reasons, assign the CA GigaStor to more than one CA Standard Monitor or CA Multi-Port Monitor. When you load-balance a CA GigaStor between two monitoring devices, the CA GigaStor Connector automatically distributes traffic to each monitoring device based on the IP address of one of the endpoints:

- The server IP address, if it is a client-server connection
- The lowest IP address, in the case of a multi-tier application, server-server connection


The CA GigaStor Connector sends packet digest files for traffic from servers with an even IP address to one monitoring device, and traffic from servers with an odd IP address to the other.

As shown in the following example, you can assign a CA GigaStor to more than one CA Standard Monitor or CA Multi-Port Monitor.



If your deployment consists of a single management console and a CA GigaStor, assign the CA GigaStor to the monitor on the management console.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
3. Scroll to the ADA Monitoring Device List and click  to edit the CA Standard Monitor or CA Multi-Port Monitor to which you want to assign the CA GigaStor. Use the Active Feeds column to determine whether a CA GigaStor is assigned to a monitor.
4. Click Monitor Devices in the third Show Me menu.
Monitor Devices opens.
5. Click a CA GigaStor from the GigaStor Devices list and click OK.
6. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.
Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

More information:

[Monitoring with a CA Standard Monitor](#) (see page 265)

[Install and Configure Software on the GigaStor Appliance](#) (see page 326)

Install CA Observer on the User's Computer

To enable a management console user to open a packet capture investigation on a CA GigaStor, the following software components must be installed on the user's computer:

- CA Observer Expert. CA Observer Expert with CA GigaStor extends global application response-time monitoring from CA Application Delivery Analysis to provide drill-down into packet-level data for root-cause analysis.

CA Observer Expert must be able to communicate with the CA GigaStor on TCP-25901 through TCP-25903. For more information about installing and licensing CA Observer Expert on the management console computer, contact your CA GigaStor administrator. Unlike the Connector for CA Observer Expert, the user cannot install CA Observer Expert from the management console.

- CA Observer Expert Connector. When a management console user attempts to open a packet capture investigation on a CA GigaStor for the first time, the Packet Capture Investigation report displays a link that prompts you to install the CA Observer Expert Connector.

Under Required Software, click the arrow link to run the Connector setup program.

Packet Capture Investigation						
Server:	191.168.2.50 (191.168.2.50)					
Date:	09/28/2009 14:12 CDT					
Investigator:	10.0.5.209 (10.0.5.209)					
Server		Application		Network		
Name	Address	Name	Port	Name	Subnet	
191.168.2.50	191.168.2.50	HTTP	80	BP1	191.168.100.0/24	
GigaStor Probe			Timeframe			
Address	Instance Name	Start Time		End Time		
10.0.1.3	Network 1	2009-09-28 14:12:09.000		2009-09-28 14:13:09.000		
Result	Download Observer Filter File			Duration	Required Software	
✓	191.168.2.50.P80.2009-09-28T191209UTC.soc			1.0 min	↓	

Give the User Permission to a Passive Probe Instance

To enable a management console user to open a packet capture investigation on a CA GigaStor in CA Observer, give the CA Application Delivery Analysis user permission to a passive probe instance on the CA GigaStor.

If the CA GigaStor has more than one active probe instance, give the user permission to a passive probe instance that corresponds to the appropriate active probe instance. You cannot give a management console user permission to more than one passive probe instance on the same CA GigaStor. By default, the management console connects to the first active probe instance on the CA GigaStor. For information about enabling the management console to connect to another active instance on a CA GigaStor, contact CA Support.


For information about probe instance administration, see the CA GigaStor product documentation.

Follow these steps:

1. Click the Administration page.
2. Click Security, and GigaStor Instances in the Show Me menu.

GigaStor Instances By User opens.

If the GigaStor Instances command is not displayed, verify that at least one CA GigaStor has been [added](#) (see page 327) to the management console.

3. Click  to edit a management console user.

GigaStor Instances for <user> opens.

4. Type the passive instance name you want the management console user to access on the CA GigaStor and click OK.

GigaStor Instances by User refreshes the GigaStor Device [Instance Name] column to display the passive instance name on the GigaStor.

Block a CA GigaStor Input Port

If you notice there are twice the number of expected observations and doubly-long NRTT, you might need to exclude all packets from selected Monitor ports on the CA GigaStor device. The port-blocking feature is an effective way to monitor in multi-tier environments. If you have configured the CA GigaStor so that different ports report the same server traffic from different vantage points, for example, from an access switch and a distribution switch, you can exclude data from the distribution switch so that those packets are not seen as duplicates of packets collected from the access switch.

Follow these steps:

1. Log in to the CA GigaStor.
2. Change directories to C:\NetQoS\GigaStorReader and then create a file named BlockedGigaStorPorts.ini.
3. Open BlockedGigaStorPorts.ini and add an entry using the following format:
/exclude port=port


Where *port* is a port value that ranges from zero to 7, which maps to the 1 to 8 range in the CA GigaStor interface. To specify more than one input port on the CA GigaStor, separate each entry with a comma.

Edit a CA GigaStor Monitoring Device

Edit a CA GigaStor monitoring device, for example, to:

- View version information for the CA GigaStor Connector and CA Observer Expert software on the CA GigaStor appliance.
- View the status of the processes on the CA GigaStor that are responsible for sending packet digest files to the assigned monitor.
- Assign an incident response to the monitoring device
- Perform basic operations on a particular CA GigaStor, such as starting and stopping the CA GigaStor Connector.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
3. Click  to edit a CA GigaStor in the GigaStor Device List.
GigaStor Properties opens.
4. Review the status information, make any changes to the specified settings, and click OK.

For information about setting CA GigaStor properties, click Help.

To view monitoring device incidents for the CA GigaStor, click Incidents in the third Show Me menu.

More information:


[View Monitoring Device Incidents](#) (see page 254)

Edit the GigaStor Monitor Feed


The *GigaStor monitor feed* receives packet digest files from a CA GigaStor monitoring device. Edit a GigaStor monitor feed to:

- Assign a particular domain. By default, a new monitor feed is assigned to the Default Domain. If you are not using domains to separate duplicate IP traffic, this is not applicable.
- Create a pair of monitor feeds to enable the management console to automatically collect data for assigned servers from both monitor feeds.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
3. Browse the Active Feeds column in the ADA Monitoring Device List to find a monitoring device with an active GigaStor monitor feed.
4. Click  to edit the monitoring device.

If you are not sure which monitoring device to edit, use the Assigned To column in the GigaStor Device List to find the monitoring device to which a CA GigaStor is assigned.

5. In the properties dialog box for the monitoring device, scroll down to Monitor Feeds.
6. Click  to edit the GigaStor monitor feed.
7. Make your changes to the GigaStor monitor feed settings, and click Update.
For information about setting GigaStor monitor feed properties, click Help.
8. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

More information:

[Managing Tenants](#) (see page 95)

[Create a Pair of Monitor Feeds](#) (see page 243)

[How Monitoring Devices Work](#) (see page 239)



Unassign a CA GigaStor

Unassign a CA GigaStor to remove it as a source of response time data. When you unassign a monitoring device, any servers that were pinned to the corresponding monitor feed are unpinned, and another monitor feed is automatically assigned. It can take up to 10 minutes to update the monitor feed assignment.

If you have pinned servers to a GigaStor monitor feed, and you want to continue monitoring server traffic while the GigaStor is temporarily unassigned, consider the following option:

- Pin the servers to another monitor feed before you unassign the GigaStor. When you reassign the GigaStor, pin the appropriate servers to the GigaStor monitor feed.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
3. Scroll to the ADA Monitoring Device List and click  to edit the CA Standard Monitor or CA Multi-Port Monitor to which the CA GigaStor is assigned.
4. Click Monitor Devices in the third Show Me menu.
5. Click  in the Assigned Device List to unassign the CA GigaStor.
6. Click Continue with Unassigning at the prompt to proceed.

The management console updates the Assigned Device List.

7. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

GigaStor Incidents


The management console automatically creates an Major monitoring device incident if a CA GigaStor:

- Stops sending data to the management console for more than 1 hour
- Does not process more than 5 percent of the packets it receives




More information:

[Edit Monitoring Device Incident Thresholds](#) (see page 255)

Perform Basic Operations

In the CA GigaStor Device list, use the blue gear menu  to perform basic operations on all of your CA GigaStor appliances, including start and stop the CA GigaStor Connector, and synchronize monitoring.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
3. Click the blue gear menu  in the GigaStor Device List to perform basic operations on all of your CA GigaStor appliances. To perform a basic operation, browse the list and click the edit icon , and click the blue gear menu .

Start

Start the Connector. While the Connector has a Running status, it sends packet digest files to its assigned monitor for management console reports.

Stop

Stop the Connector. While the Connector has a Stopped status, the CA GigaStor continues to write packets to disk, but the Connector does not send packet digest files to its assigned monitor. Data gaps in management console reports can be resolved by viewing the actual data on the CA GigaStor using the CA Observer Expert.

Synchronize Monitor Devices

Synchronize the CA GigaStor to collect performance data based on the client networks, server subnets, and applications that are currently defined on the Console.

More information:

[How Monitoring Device Synchronization Works](#) (see page 242)

Delete a CA GigaStor Monitoring Device

Delete a CA GigaStor monitoring device to remove it as a source of response time data for the management console. If you no longer want to use a CA GigaStor as a monitoring device, unassign the CA GigaStor from its assigned monitor, and then delete the CA GigaStor.

Alternatively, you can [assign a CA GigaStor](#) (see page 328) to another CA Standard Monitor or CA Multi-Port Monitor.

More information:


[Assign a CA GigaStor to a Monitoring Device](#) (see page 328)

Delete a CA GigaStor

Delete a CA GigaStor to remove it as a source of response time data for the management console.

To delete a CA GigaStor, it cannot be assigned to a CA Standard Monitor or CA Multi-Port Monitor. After you unassign a CA GigaStor, [delete the CA GigaStor](#) (see page 335) from the management console.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
3. Click  to delete a CA GigaStor from the GigaStor Device List.
4. Click Continue with Delete at the prompt to delete the CA GigaStor.

The CA GigaStor is removed from the GigaStor Device List.

5. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

Troubleshoot a CA GigaStor Monitoring Device

Troubleshoot the CA GigaStor to identify the cause of missing report data that should have been monitored by the CA GigaStor. After you add a CA GigaStor monitoring device, it can take up to 10 minutes for the management console to report performance data.

View Active Sessions on the GigaStor Monitor Feed

Use the Active Sessions page to report on the number of active IPv4-based TCP sessions reported by the GigaStor monitor feed during the last 5-minute reporting interval. The GigaStor monitor feed includes session information from any CA GigaStor appliances that are assigned to the monitoring device.

[View active sessions information](#) (see page 245) to verify that a monitor feed is monitoring TCP sessions. The management console reports the number of active TCP sessions on a server by application port. If the monitor feed does not have any active sessions for a server or application, you have a CA GigaStor configuration problem.

View GigaStor Counter Statistics

The GigaStor counter displays information about the packet digests that it receives from its assigned CA GigaStor Connector, including information about the Netflow packets that transport the packet digests.

To view the GigaStor counter, [log into the CA Standard Monitor](#) (see page 287) where the CA GigaStor is assigned.

Important: Before you begin, synchronize the monitoring devices. The counter windows do not display until after you have synchronized data monitoring.

The GigaStor counter displays the following statistics:

Good Flows

Indicates the number of Netflow packets that were received in the order they were sent.

Dropped Flows

Indicates the number of Netflow packets that were not processed by the CA ADA Monitor service. Response time data from any packet digests that were included in the dropped Netflow packets is not included in the management console reporting.

Out of Order Flows

Indicates the number of Netflow packets that were received by the monitor, but were not received in the order they were sent.

To Server Packets

Indicates the number of packets that were sent from a client to a server.

From Server Packets

Indicates the number of packets that were sent from a server to a client.

To Server Bytes

Indicates the number of bytes that were sent from a client to a server.

From Server Bytes

Indicates the number of bytes that were sent from a server to a client.

Total Seen Packets

Indicates the number of packets that were inspected by the CA GigaStor Connector to determine whether the packet header matched a specified application port, client network, and server subnet.

Total Captured Bytes

Indicates the total byte count for packets that match a specified application port, client network, and server subnet.

Note The monitor inspects each packet header to determine whether the packet matches a specified application port, client network, and server subnet. For more information, see Total Seen Packets.

Accepted Sessions

Indicates the number of TCP sessions that match a valid application/server/network combination on the management console.

Rejected for Server

Indicates the server IP did not match a server subnet monitored by the management console.

Rejected for Client

Indicates the client IP did not match the list of client networks to be monitored by the management console.

Rejected for Port

Indicates the server port matched the list of ports to be ignored by the management console.

Rejected for Positive

Reserved for future use.

More information:

[How Client Networks Work](#) (see page 29)

[How Applications Work](#) (see page 103)

[How Servers Work](#) (see page 67)

Chapter 16: Monitoring with Cisco WAAS

This section contains the following topics:

[How Cisco WAAS Works as a Monitoring Device](#) (see page 342)

[Add a Cisco WAE Monitoring Device](#) (see page 349)

[Edit a Cisco WAE Monitoring Device](#) (see page 352)

[Edit the WAN Opt Monitor Feed](#) (see page 353)

[Unassign a Cisco WAE](#) (see page 354)

[WAAS Incidents](#) (see page 354)

[Delete a Cisco WAE Monitoring Device](#) (see page 355)

[Disable Flow Monitoring on a Cisco WAE](#) (see page 355)

[Reset Optimized Applications](#) (see page 356)

[Troubleshoot a Cisco WAE Monitoring Device](#) (see page 357)

[Monitor a Server with a Group of Cisco WAE Devices](#) (see page 361)

[Share Optimization Data between Management Consoles](#) (see page 365)


How Cisco WAAS Works as a Monitoring Device

A Cisco WAE device (Cisco WAE) serves as a type of monitoring device for CA Application Delivery Analysis. Use Cisco WAE devices to gain visibility into optimization. Unlike a monitoring device that monitors the server SPAN, Cisco WAE devices are distributed across the network.

Monitoring WAN optimization between Cisco WAE devices lets you see how Cisco WAAS optimization affects individual application response times at each segment of the network.

In the example below, the Cisco WAE devices at the branch and data center locations send optimized application performance data to a CA Standard Monitor that also monitors the server SPAN. The CA Standard Monitor:

- Calculates performance metrics for the optimized traffic on the Client and WAN segments.
- Replaces the Server segment performance data from the data center WAE with the more accurate performance data collected by the CA Standard Monitor from the server's mirrored switch port.
- Automatically monitors application traffic in the server SPAN and updates the management console with a list of servers to be monitored by all monitoring devices.



To monitor a Cisco WAAS environment, you are not required to monitor the server's mirrored switch port.

More information:

[Monitoring Device Considerations](#) (see page 348)

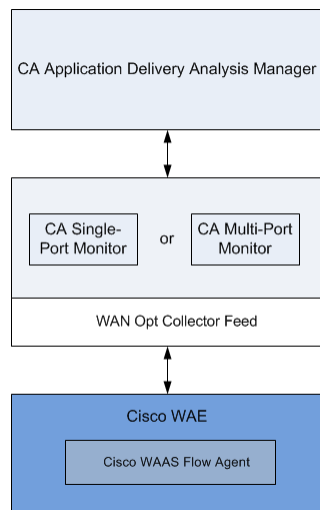
How Cisco WAAS Works

The Cisco WAAS Flow Agent on the Cisco WAE device polls the management console every 5 minutes for a list of server IP addresses to monitor. The Cisco WAAS Flow Agent sends packet digest files with matching TCP headers of optimized traffic to its assigned CA Standard Monitor or CA Multi-Port Monitor. A Distributed management console is required.

The packet digest files are received by the WAN Opt monitor feed on the CA Standard Monitor or CA Multi-Port Monitor where they are processed into response time metrics for each observed application/server/network combination.

The Cisco WAAS Flow Agent monitors application traffic based on the list of servers stored on the management console. To enable the Cisco WAAS Flow Agent to automatically monitor new server traffic, an additional monitoring device is required to monitor the server SPAN.

The Cisco WAAS Flow Agent does not send TCP headers for unoptimized traffic and does not calculate response time metrics.



Unlike the Cisco NAM Metric Agent, the Cisco WAAS Flow Agent sends:

- Packet digest files that include TCP headers rather than computed response time metrics.
- Sends TCP headers for optimized traffic that matches the Server List rather than all optimized traffic.

How Monitor Feed Assignment Works

The management console automatically combines response time metrics from all Cisco WAE devices to calculate application response time on each network segment. If there is a monitor that is monitoring the server SPAN, and the monitor is closest to the server, the management console assigns that monitor feed to monitor the Server segment. Otherwise, the management console assigns the WAN Opt monitor feed to a server.

How Network Segments Work

Because the management console is now monitoring from multiple points on the network for a single application-server-network combination, the management console generates a separate set of metrics for each of the three network segments and treats each network segment as a separate application. The management console generates a separate set of metrics for the:

- Client segment, which is the network segment between the client IPs in the branch location and the branch WAE device. To report this network segment, the management console requires the branch WAE device to export response time data.
- WAN segment, which is the network segment between the branch WAE device and the data center WAE device. To report this network segment, the management console requires the data center WAE device to export response time data.
- Server segment, which is the network segment between the data center WAE device and the data center servers. To report this network segment, the management console requires the data center WAE device to export response time data.

When using a monitoring device to monitor the server SPAN, the management console replaces the Server segment data from the data center Cisco WAE with the more accurate server SPAN data. It is more accurate because the server's SPAN source is closer to the actual server. If the management console sees unoptimized application traffic in the server SPAN, it generates a separate set of metrics for the unoptimized application, for example, SMTP.

The management console appends the network segment to the application name, for example, SMTP [Client], SMTP [WAN], and SMTP [Server]. In the example below, the SMTP application represents the unoptimized application performance. Note that the unoptimized application response time is faster because the data comes from local users in the data center.

Performance by Application				
Application	Port(s)	Transaction Time		Observations
		Wtd. Average: 7.12 ms		Average: 69.29 ms
FTP [WAN]	21		85.17 ms	6,729,198
FTP [Server]	21		82.20 ms	4,746,399
FTP [Client]	21		68.76 ms	4,765,733
FTP	21		59.94 ms	1,206
SMTP [Server]	25		39.78 ms	981,705
SMTP [Client]	25		28.26 ms	995,452
SMTP [WAN]	25		26.07 ms	1,747,521
SMTP	25		9.10 ms	10,251,044

How Performance Thresholds Work for WAN-Optimized Network Segments

The management console creates separate applications to report on application performance across the Client, WAN, and Server segments of the network. Customize the application performance thresholds on each network segment to make the management console more or less sensitive to variations in performance.

More information:

[Edit Performance Thresholds for WAN-Optimized Network Segments](#) (see page 156)

How Monitoring Works when Optimization Stops

Depending on how long optimization is interrupted, the management console uses different reporting methods for the affected traffic. If necessary, you can reset the management console to disregard recent changes in optimization and report optimized and unoptimized traffic based on the optimization information that is currently available. For more information, see the following sections.

Temporary Interruptions

If the management console temporarily stops receiving segmented application performance data from a Cisco WAE, the management console uses data from the unoptimized server SPAN to report on the application in the Server segment. The management console uses several criteria to determine whether an interruption is temporary or permanent, but a temporary interruption is typically less than 20 minutes.

While the management console temporarily reports the unoptimized application data from the server's mirrored switch port in the Server segment:

- All metrics on the Optimization page are accurate, although not all metrics will populate.
- Measurements for the unoptimized sessions do not affect the Client or WAN segments.
- The Server segment is accurate because the management console always uses the more accurate data from the server's mirrored switch port.
- Network Round Trip Time [Server], which is displayed on the Engineering page, shows an increase because it no longer gets 100% local ACKs. The spillover measurements show actual Network Round Trip Time out to the client.
- RetransDelay [Server] and PacketLossPct [Server] might also show increases.

When optimization data resumes, the management console automatically reports the segmented application data in the application's Client, WAN, and Server segments.

Temporary interruptions to optimized monitoring can occur, for example, when:

- Optimization stops. This can occur when a Cisco WAE does not have the resources to optimize additional sessions. The unoptimized sessions are called *spillover*.
- Monitoring stops. This can occur when a CA Standard Monitor or CA Multi-Port Monitor stops receiving packet digests from its assigned Cisco WAE devices, for example, when there is a disruption in the link between a Cisco WAE and its assigned monitoring device, or when the Cisco WAE is not configured to export response time data.

Permanent Changes

The management console uses several criteria to determine whether an interruption is temporary or permanent, but when a Cisco WAE stops sending segmented application performance data for more than 20 minutes, the management console typically considers the interruption to be permanent. In this case, the management console stops reporting the unoptimized traffic from the server's mirrored switch port in the application's Server segment, for example, HTTP [Server], and instead reports on the SPAN data in the unoptimized application, HTTP. If optimization resumes, the management console automatically resumes monitoring the optimized application by network segment.

If you are setting up WAAS optimization, for example, to measure the benefits of optimizing different applications, and you want the management console to immediately report changes in optimization, you can [reset](#) (see page 347, see page 356) the management console.

Sizing Recommendations

A Cisco WAE device sends packet digest files to the Management NIC on its assigned CA Standard Monitor or CA Multi-Port Monitor for processing.

A CA Standard Monitor or CA Multi-Port Monitor can process packet digest files from all three segments (Client, WAN, and Server) for at least 50,000 optimized connections. If possible, avoid assigning more than one data center WAE to the same monitor. Load-balance the branch Cisco WAE devices among the available CA Standard Monitor or CA Multi-Port Monitors.

To avoid overloading the Management NIC on a CA Standard Monitor or CA Multi-Port Monitor, avoid assigning Cisco WAE devices to a monitor that is also receiving packet digests from a CA GigaStor.

Monitoring Device Considerations

When using a Cisco WAE as a monitoring device, keep in mind:

- The management console requires the branch WAE device to export application response times on the Client network segment. If necessary, you can simply enable the data center WAE device to report application performance to the management console across the WAN and Server segments.
- The management console cannot report application performance by network segment for:
 - A FTP application because Cisco WAAS does not optimize FTP.
 - A Web application by URL. Alternatively, you can define a Standard application to monitor all TCP-80 traffic across the Client, WAN, and Server segments.
- If communication is interrupted between a Cisco WAE device and its assigned CA Standard Monitor or CA Multi-Port Monitor, the Cisco WAAS Flow Agent temporarily stores its packet digest files to avoid the loss of reporting data.
- The Cisco WAAS Flow Agent polls the management console every 5 minutes for a list of server IP addresses to monitor.

To enable a Cisco WAE device to monitor new server traffic that matches a server subnet, configure a monitoring device to monitor the server SPAN, such as a CA Multi-Port Monitor. If you are only using Cisco WAE devices to collect response time data, manually update the Console to add the server IP addresses that you want to monitor.

To view the list of servers monitored by the management console, click Data Monitoring, Servers in the Show Me list.

- A Cisco WAE device does not capture performance information about unoptimized, pass-through traffic. To monitor application performance for pass-through applications, configure a monitoring device to monitor the server SPAN, such as a CA Multi-Port Monitor.
- A Cisco WAE device recognizes and eliminates duplicate data, therefore, it is unnecessary to configure a CA Standard Monitor or CA Multi-Port Monitor to de-duplicate packets from the WAN Opt monitor feed.
- Packet capture investigations cannot be performed from a Cisco WAE device, however, if you monitor the server SPAN, for example, with CA Multi-Port Monitor, the management console uses this device to take the packet capture.

Add a Cisco WAE Monitoring Device

To add a Cisco WAE monitoring device to the management console, perform the following tasks:

1. [Configure the Cisco WAE](#) (see page 350) device to export application response time data.
2. [Assign the Cisco WAE](#) (see page 351) to a CA Standard Monitor or CA Multi-Port Monitor.
3. Wait up to 10 minutes for the Cisco WAE device to poll the CA Application Delivery Analysis Manager for a list of servers to monitor and for the management console to display data for an optimized application.

If necessary, [troubleshoot the Cisco WAE](#) (see page 357) device to confirm that the Cisco WAE device is monitoring the optimized applications.

Prerequisites

To add a Cisco WAE monitoring device, you must meet the following prerequisites:

- The Cisco WAE is running a supported version of the Cisco WAAS software. The management console supports Cisco WAAS 4.0.17 to 4.4.3a, sending packet digest files to a:
 - CA Standard Monitor
 - CA Multi-Port Monitor
- The time and date are the same on each Cisco WAE and on the assigned CA Standard Monitor or CA Multi-Port Monitor. The performance data collected by a Cisco WAE is time stamped using the device's timestamp. If you configured the management console to use the Network Time Protocol (NTP), use the WAAS Central Manager to configure NTP on the Cisco WAE devices. For more information, see the *Cisco Wide Area Application Services Configuration Guide*.
- The [monitoring device ratio](#) (see page 249) is sized properly.
- The Cisco WAE can communicate with its assigned CA Standard Monitor or CA Multi-Port Monitor, and the management console, on TCP-7878.

Configure the Cisco WAE to Export Response Time Data

From the Cisco WAE CLI or Central Manager GUI, configure the Cisco WAE device to export application response time data.

Follow these steps:

1. On the Cisco WAE, change to configuration mode by running the following command:
config
2. The command line prompt changes to:
WAE<config>#
3. Disable flow monitoring on the Cisco WAE by running the following command:
no flow monitor tcpstat-v1 enable
4. Register the Cisco WAE with the IP address of the management console by running the following command:
flow monitor tcpstat-v1 host <MCAddress>

Where <MCAddress> is the IP address of the management console.

5. Enable flow monitoring on the Cisco WAE by running the following command:
flow monitor tcpstat-v1 enable
6. Return to privilege mode by running the following command:
exit
7. The command line prompt changes to:
WAE#
8. To confirm the Cisco WAAS Flow Agent is connected to the CA Application Delivery Analysis Manager, run the following command:
show statistics flow monitor tcpstat-v1

The results should indicate that the Configured Host Address is the IP address of the CA Application Delivery Analysis Manager. Note that it is normal for the Host Connection State to be Waiting to Poll except for when the Cisco WAE is polling the CA Application Delivery Analysis Manager.

9. Verify that the Cisco WAE appears in the list of WAN-optimization devices on the management console:
 - a. Open the management console.
 - b. Click the Administration page.
 - c. Click Data Monitoring, Monitoring Devices in the Show Me menu.
 - d. Scroll to the WAN Optimization Device List. If the Cisco WAE is displayed, you can [assign](#) (see page 351) the Cisco WAE to a CA Standard Monitor or CA Multi-Port Monitor.

Assign a Cisco WAE to a Monitoring Device

After you configure the Cisco WAE device to export response time data, you are ready to assign the Cisco WAE to a CA Standard Monitor or CA Multi-Port Monitor. Do not assign a Cisco WAE to a CA Virtual Systems Monitor.

To maximize the available monitoring device resources, disable packet monitoring when you [add the CA Standard Monitor](#) (see page 272).

When deciding where to assign a Cisco WAE, simply load-balance between CA Standard Monitor and CA Multi-Port Monitor monitoring devices. If possible, avoid assigning more than one data center WAE to the same monitor and then load-balance the branch Cisco WAE devices.


The Cisco WAE:

- Polls its assigned monitoring device every 5 minutes for the list of servers to monitor
- Sends packet digest files to its assigned monitoring device

Assign a Cisco WAE to a monitoring device that is monitoring the server SPAN or to a monitoring device that is dedicated to receiving packet digest files from Cisco WAE devices. To avoid overloading the Management Port on a CA Standard Monitor, do not assign a Cisco WAE to a CA Standard Monitor that is also receiving packet digest files from a CA GigaStor.

If you are using domains to separate duplicate IP traffic, [edit the WAN Opt monitor feed](#) (see page 353) and assign it to a domain.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
3. In the ADA Monitoring Device List, click  to edit a CA Standard Monitor or CA Multi-Port Monitor.

Note that in the Active Feeds column, WAN Opt indicates that a WAN optimization device, such as a Cisco WAE device, is assigned to a monitor.

4. Click Monitor Devices in the third Show Me menu.
5. In Monitor Devices, scroll to WAN Optimization and click a device in the Available column. Click the right-arrow to move it into the Assigned column. The Cisco WAE devices in the Assigned column are currently assigned to the monitor.

If the Cisco WAE you want is not listed, verify that the Cisco WAE is [configured to communicate with the management console](#) (see page 350).

6. Repeat to assign another Cisco WAE, or click OK to finish.

7. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

8. After 5-10 minutes, the Optimization page of the management console should display application performance data for the optimized network segments. In the Optimization page of the management console, filter by Last Hour to see the performance data as it is received.

If you do not see data for your application, [troubleshoot the Cisco WAE monitoring device](#) (see page 357).

More information:

[Add a CA Standard Monitor](#) (see page 272)

[Edit the WAN Opt Monitor Feed](#) (see page 353)

Edit a Cisco WAE Monitoring Device

Edit a Cisco WAE monitoring device to update its properties, for example, to assign an incident response. While you are editing a Cisco WAE, you can also view any monitoring device incidents.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
3. Scroll to the WAN Optimization Device List.
4. (Optional) Click a source set from the Select Source Set list to view the corresponding group of Cisco WAE devices.

5. Click  to edit a Cisco WAE monitoring device.

WAN Optimization Properties opens.

6. Click Incident Response to select an incident response and then click OK.

To view incidents for the monitoring device, click Incidents in the third Show Me menu.

More information:

[Monitor a Server with a Group of Cisco WAE Devices](#) (see page 361)


[View Monitoring Device Incidents](#) (see page 254)

Edit the WAN Opt Monitor Feed

Edit the WAN Opt monitor feed to assign a particular domain to a group of Cisco WAE devices. By default, a new monitor feed is assigned to the Default Domain. The *WAN Opt* monitor feed, which is available on a CA Standard Monitor or CA Multi-Port Monitor, receives packet digest files from WAN-optimization devices, such as a Cisco WAE device.

Note: In a WAAS environment, it is unnecessary to create a pair of WAN Opt monitor feeds. If a location has an additional Cisco WAE for redundancy, simply assign it to a CA Standard Monitor or CA Multi-Port Monitor. To enable support for redundancy in a WAAS environment, [create a pair of monitor feeds](#) (see page 243) that monitor the server SPAN.


Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
3. Scroll to the ADA Monitoring Device List and click  to edit a CA Standard Monitor or CA Multi-Port Monitor.

Note that in the Active Feeds column, WAN Opt indicates that a WAN optimization device, such as a Cisco WAE device, is assigned to a monitor.

If you are not sure which monitor to edit, use the Assigned To column in the WAN Optimization Device List to find its assigned monitor.

Monitor Properties opens.

4. Scroll to the Monitor Feeds section and click  to edit the WAN Opt monitor feed, and specify the WAN Opt monitor feed settings.

Note that the monitor feed settings for a WAN Opt monitor feed do not let you configure a Secondary Feed. A secondary monitor feed is not applicable for WAN optimization devices.

5. Click Update.
6. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.


More information:

[Managing Tenants](#) (see page 95)

Unassign a Cisco WAE

Unassign a Cisco WAE from its assigned monitor. If you unassign all WAN optimization devices from a monitor, the management console automatically removes its WAN Opt monitor feed.


Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
3. Scroll to the ADA Monitoring Device List and click  to edit the CA Standard Monitor or CA Multi-Port Monitor to which the Cisco WAE is assigned.

Note that in the Active Feeds column, WAN Opt indicates that a WAN optimization device, such as a Cisco WAE device, is assigned to a monitor.

4. Click Monitoring Devices in the third Show Me menu.

Monitor Devices opens.

5. Click  to unassign the Cisco WAE.
6. Click Continue with Unassigning at the prompt to unassign the Cisco WAE.

The management console updates the Assigned Device List.

7. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

WAAS Incidents

The management console automatically creates a Major monitoring device incident if a Cisco WAE device stops sending data to the management console for more than 15 minutes.

For more information:

[Edit Monitoring Device Incident Thresholds](#) (see page 255)

Delete a Cisco WAE Monitoring Device

Delete a Cisco WAE to remove it as a source of response time data for the management console. When you delete a Cisco WAE monitoring device, the management console automatically assigns another monitor feed to any servers where the WAN Opt monitor feed was assigned. It can take up to 10 minutes for the management console to update the monitor feed assignment.

Deleting a Cisco WAE monitoring device automatically unassigns the Cisco WAE from its CA Standard Monitor or CA Multi-Port Monitor.

Follow these steps:

1. [Unassign the Cisco WAE](#) (see page 354) from its assigned CA Standard Monitor or CA Multi-Port Monitor monitoring device. Alternatively, you can [reassign](#) (see page 351) a Cisco WAE device to another monitoring device.
2. [Disable flow monitoring on the Cisco WAE](#) (see page 355) to remove the Cisco WAE from the list of unassigned Cisco WAE monitoring devices in the management console.
3. [Delete the Cisco WAE](#) (see page 356) from the management console.

Disable Flow Monitoring on a Cisco WAE

To remove a Cisco WAE from the list of unassigned Cisco WAE monitoring devices, you must first disable flow monitoring on the Cisco WAE.


Follow these steps:

1. On the Cisco WAE, change to configuration mode by running the following command:
`config`
2. The command line prompt changes to:
`WAE<config>#`
3. Disable flow monitoring on the Cisco WAE by running the following command:
`no flow monitor tcpstat-v1 enable`
4. Return to privilege mode by running the following command:
`exit`
5. The command line prompt changes to:
`WAE#`
6. To confirm that flow monitoring is disabled, run the following command:
`show statistics flow monitor tcpstat-v1`
7. The results of the command confirm flow monitoring is disabled:
Flow application is not enabled or is not available.

Delete a Cisco WAE Monitoring Device

Before you delete a Cisco WAE monitoring device, configure the Cisco WAE to disable the export of response time data. If you delete a Cisco WAE monitoring device from the management console, and the Cisco WAE device is configured to export response time data, the Cisco WAE will continue to be available as a monitoring device.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
3. Scroll to the WAN Optimization Device List and click  to delete a Cisco WAE monitoring device.
4. Click Continue with Delete at the prompt to delete the Cisco WAE.

The Cisco WAE is removed from the WAN Optimization Device List.


Reset Optimized Applications

Reset all optimized applications to enable the management console to report application performance based on the segmented application data that the management console currently receives. [Normally](#) (see page 345), you do not need to reset optimization because, if that optimization is interrupted, the management console continues to report on the application.

If you are planning to use the management console to prove the benefits of optimizing different applications, resetting optimization lets you quickly compare changes in WAAS optimization. For example, if you stop optimization, the management console continues to report the unoptimized server SPAN in the application's Server segment for up to 20 minutes before reporting the unoptimized traffic in the unoptimized application.

After you reset optimization, any unoptimized application traffic that was reported in Server segment is now reported in the unoptimized application, and currently optimized applications are reported by network segment. However, it can take up to 10 minutes for the management console to substitute the server SPAN metrics into the application's Server segment.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
3. Scroll to the WAN Optimization Device List, click  and select Reset Optimized.
4. Wait 10 minutes to see updated report data in the Optimization page.

More information:

[How Monitoring Works when Optimization Stops](#) (see page 345)

Troubleshoot a Cisco WAE Monitoring Device

Troubleshoot the Cisco WAE to identify the cause of missing report data. After you add a Cisco WAE monitoring device, it can take up to 10 minutes for the management console to report optimized application performance by network segment.

If the management console does not display segmented application data for a branch location or all branch locations, confirm the corresponding Cisco WAE device has active sessions, and the Cisco WAE device is exporting response time data for the application traffic you want.

If network segment data is missing for the:

Client segment

Verify the branch WAE device.

WAN segment

Verify the data center WAE device.

Server segment

Verify the monitoring device that monitors the traffic between the data center WAE device and the data center servers. If a monitoring device is not configured to monitor the server SPAN, verify the data center WAE device.

View Active Sessions

Use the Active Sessions page to report on the number of active sessions reported by the WAN Opt monitor during the last 5-minute reporting interval.

Use the active sessions information to verify that the WAN Opt monitor feed is [monitoring TCP sessions](#) (see page 245). The management console reports the number of active TCP sessions on a server by application port. If the monitor feed does not have any active sessions for a server or application, you have a WAE configuration problem.

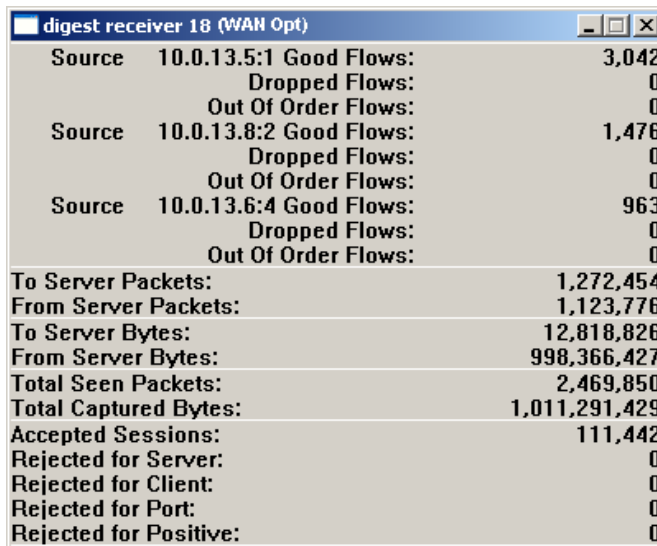
View WAN Opt Counter Statistics

View the WAN Opt counter to display information about the packet digests that it receives from assigned Cisco WAE devices, including information about the Netflow packets that transport the packet digests.

To view the WAN Opt counter window, you must [log into](#) (see page 287) the CA Standard Monitor to which the Cisco WAE is assigned.

Important: Before you begin, synchronize the monitoring devices. The counter windows do not display until after you have synchronized data monitoring.

The WAN Opt counter displays the statistics from the Cisco WAE devices that are assigned to the monitor:



digest receiver 18 (WAN Opt)	
Source 10.0.13.5:1	Good Flows: 3,042
	Dropped Flows: 0
	Out Of Order Flows: 0
Source 10.0.13.8:2	Good Flows: 1,476
	Dropped Flows: 0
	Out Of Order Flows: 0
Source 10.0.13.6:4	Good Flows: 963
	Dropped Flows: 0
	Out Of Order Flows: 0
To Server Packets:	1,272,454
From Server Packets:	1,123,776
To Server Bytes:	12,818,826
From Server Bytes:	998,366,427
Total Seen Packets:	2,469,850
Total Captured Bytes:	1,011,291,429
Accepted Sessions:	111,442
Rejected for Server:	0
Rejected for Client:	0
Rejected for Port:	0
Rejected for Positive:	0

Good Flows

Identifies the number of Netflow packets that were received in the order they were sent.

Dropped Flows

Identifies the number of Netflow packets that were not processed by the CA ADA Monitor service. Response time data from any packet digests that were included in the dropped Netflow packets is not included in the management console reporting.

Out of Order Flows

Identifies the number of Netflow packets that were received by the monitor, but were not received in the order they were sent.

To Server Packets

Identifies the number of packets that were sent from a client to a server.

From Server Packets

Identifies the number of packets that were sent from a server to a client.

To Server Bytes

Identifies the number of bytes that were sent from a client to a server.

From Server Bytes

Identifies the number of bytes that were sent from a server to a client.

Total Seen Packets

Identifies the number of packets that were inspected to determine whether the packet header matched a specified application port, client network, and server subnet.

Total Captured Bytes

Identifies the total byte count for packets that match a specified application port, client network, and server subnet.

Note The Cisco WAAS Flow Agent inspects each packet header to determine whether the packet matches a specified application port, client network, and server subnet. For more information, see Total Seen Packets.

Accepted Sessions

Identifies the number of TCP sessions that match a valid application/server/network combination in the management console.

Rejected for Server

Identifies the server IP did not match a server subnet monitored by the management console.

Rejected for Client

Identifies the client IP did not match the list of client networks to be monitored by the management console.

Rejected for Port

Identifies the server port matched the list of ports to be ignored by the management console.

Rejected for Positive

Reserved for future use.

More information:

[How Client Networks Work](#) (see page 29)

[How Applications Work](#) (see page 103)

[How Servers Work](#) (see page 67)

Verify the Cisco WAE Configuration

Verify the Cisco WAE is:

- Optimizing the servers you want to monitor.
- The Cisco WAAS Flow Agent is exporting response time data for the optimized servers.

Follow these steps:

1. On the Cisco WAE, run the following command to show its optimized traffic:

WAAS/WAE version 4.1.x

```
show statistics connection all
```

WAAS/WAE version 4.0.x

```
show tfo connection summary
```

2. Review the list of server IP addresses and ports to make sure the Cisco WAE is optimizing the application traffic of interest.
3. Run the following command to view the list of servers for which Cisco WAAS Flow Agent is exporting response time data:

```
show statistics flow filters
```

If the list of servers does not match the list of optimized servers from the previous step:

- a. Verify the management console is monitoring the server.
 - b. Verify the server is assigned to the application.
4. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.
 5. Wait 5 minutes and then verify that the list of servers which the Cisco WAE is monitoring corresponds to the list of servers on the management console by running the following command:

```
show statistics flow filters
```

You should see that the server has now been added and the number of flow hits is greater than 0.

More information:

[Manage Servers](#) (see page 76)

[Assign Servers to an Application](#) (see page 126)

Monitor a Server with a Group of Cisco WAE Devices

It is generally a good idea to allow all Cisco WAE devices to monitor all servers because Cisco WAE devices are distributed throughout a network, and individual Cisco WAE devices might see the following:

- All the traffic to and from a specific server
- None of the traffic to and from a specific server
- Some of the traffic to and from a specific server

How Source Sets Work

A *source set* is a group of Cisco WAE devices. Most of the time, you will not need to configure source sets, as the management console assigns all servers and Cisco WAE devices to the same source set so that all Cisco WAE devices monitor all servers. If necessary, you can create a source set to specify a group of Cisco devices that you want to monitor a particular server.

A source set does not ensure the uniqueness of the traffic reported by those devices. If necessary, [use domains to separate duplicate traffic](#) (see page 95).

To monitor a server with a group of Cisco WAE devices, perform the following tasks:

- [Assign a source set to the appropriate Cisco WAE Devices](#) (see page 362).
- [Assign the source set to the appropriate servers](#) (see page 363).

Assign a Source Set to a Cisco WAE Device


Assign a source set to a Cisco WAE device to create a group of Cisco WAE devices. For reporting purposes, assign the source set to a server to ensure that only Cisco WAE devices in the source set monitor WAN optimized application traffic on the server.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
3. Scroll to the WAN Optimization Device List and click Select Source Set to filter the list of WAE devices.

If the Select Source Set list is not displayed, no additional source sets are defined.

By default, the management console assigns a new WAE device to the default source set, which is named Default Deployment. Choose Default Deployment to view the list of WAE devices that are currently assigned to the default source set.

4. Click  to edit a Cisco WAE monitoring device.
WAN Optimization Device Properties opens.
5. Click Select Source Set to choose the source set to which you want to assign the Cisco WAE or click Add to add a new source set, type a source set name, and click OK.

The management console updates the list of Cisco WAE devices that belong to the currently selected source set.

If the Source Set list is not displayed, no additional source sets are defined.

6. Repeat these steps to assign a source set to each Cisco WAE monitoring device.


Assign a Source Set to a Server

Most of the time, you will not need to assign a source set to a server, as the management console automatically monitors all servers with all Cisco WAE devices.

If necessary, specify a [group of Cisco WAE devices](#) (see page 362) to monitor a server by assigning a source set to the server.

If you create one or more source sets, the Server Properties page displays a Source Set option that lets you assign a particular source to a server.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
3. Scroll to the Servers List and click  to edit a server.

Server Properties opens.


4. Click Source Set to select a different source set and click OK.

If the Source Set list is not displayed, no additional source sets are defined.

Rename a Source Set

Edit a source set to change its name. If you want to change the source set to which a Cisco WAE monitoring device is assigned, [edit the Cisco WAE monitoring device](#) (see page 352).


Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
3. Scroll to the WAN Optimization Device List and click Select Source Set to filter the list of WAE devices.
4. Click  and click Edit Source Set.
5. In Source Set, type a new name for the source set and click OK.

Delete a Source Set

When you delete a source set, the management console reassigns the corresponding Cisco WAE devices and any servers that were assigned to the deleted source set to the Default Deployment source set.

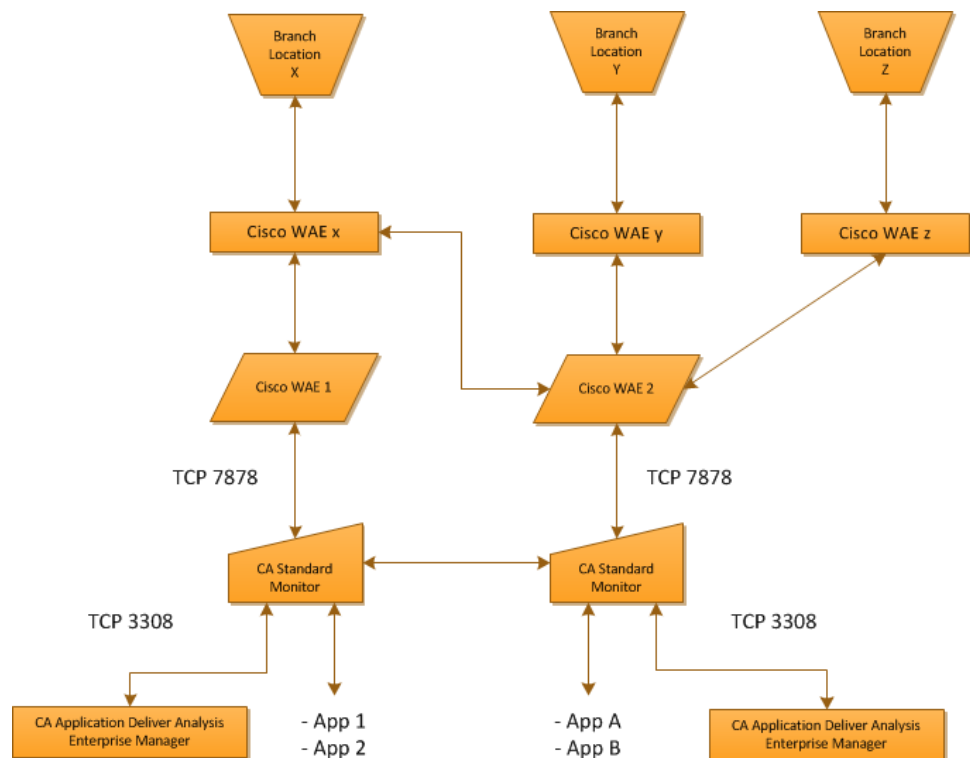
Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
3. Scroll to the WAN Optimization Device List and click Select Source Set to choose the source set you want to delete.
4. Click  and click Delete Source Set.
5. In Delete Source Set, you are prompted to confirm the delete.

Share Optimization Data between Management Consoles

A management console can typically support all the Cisco WAE devices in your environment. However, if your WAAS deployment optimizes an application that requires you to deploy more Cisco WAE monitoring devices than can be supported by a single management console, share the Client, WAN, and Server segment application performance data for your remote sites between more than one management console.

In the example below, sharing WAN-optimized application performance data enables the management console on the right to report the Client segment for App A and App B segment traffic between Branch X and WAE x. Without sharing, the management console on the right would only report the WAN and Server segments for App A and App B on Branch X.



Alternatively, if you configure the management console on the left to monitor the performance of App A and App B across your remote branches, sharing WAN-optimized application performance data enables the console to report the WAN and Server segments for App A and App B on Branch X, plus the Client, WAN and Server segments for App A and App B on Branch Y and Branch Z. Without sharing, the management console on the left would not report the WAN and Server segments for App A and App B on Branch X, or Branch Y and Branch Z.

The management console shares packet digest files from Cisco WAE devices between more than one management console by polling each management console for its list of CA Standard Monitor and CA Multi-Port Monitor monitoring devices, and then sharing packet digest files with a monitor that belongs to the shared Console. For example, if a CA Standard Monitor receives packet digest files for a server that it is not supposed to monitor, the monitor shares the packet digest files with a monitor that belongs to a shared management console. To minimize the amount of data that is shared, monitor an application from the management console that owns the monitor that is closest to the application.

Share WAN-Optimized Performance Data

Sharing WAN-optimized application performance data requires you to:

- Assign all the Cisco WAE devices and servers to the same source set. If you have not configured a source set, you can simply use the Default Deployment source set.
- Configure the monitoring devices that receive packet digest files to share application performance data.
- Ensure that all the shared monitoring devices can communicate with:
 - Each management console on TCP-3308. All the monitoring devices that receive packet digest files from a Cisco WAE must be able to communicate with each management console on TCP-3308.
 - All the Cisco WAE devices on TCP-7878. The management console automatically configures sharing between monitoring devices.

After you enable sharing on all the monitoring devices, it can take up to 25 minutes for a Cisco WAE to poll its assigned monitor for an updated list of servers and to begin reporting shared data.

If a monitor cannot share WAN-optimized application performance data, it will store the data until sharing is restored. However, a monitor can only store up to 1 GB of shared data. Any shared data that cannot be stored on the monitor is lost and cannot be recovered.

Follow these steps:

1. Ensure that the CA Standard Monitor and CA Multi-Port Monitor monitoring devices which receive packet digest files from a Cisco WAE can communicate with each management console on TCP-3308. All the monitoring devices must be able to communicate with management console on TCP-3308.
2. Ensure that all monitoring devices can communicate with each other on TCP-7878.
3. Configure all monitoring devices to share WAN-optimized application performance data:
 - a. Create a configuration file named DTMDistributedConsoles.ini.
 - b. In DTMDistributedConsoles.ini, enter the IP address of each management console you want to share. Be sure to include the IP address of the management console that is assigned to the monitor. When specifying IP addresses, separate each IP address on a new line using dot-decimal notation.
 - c. Copy the same DTMDistributedConsoles.ini file to all monitoring devices. For a:
 - CA Standard Monitor, copy the file to the <ADA_HOME>\bin folder
 - CA Multi-Port Monitor, copy the file to the /opt/NetQoS/bin folder
 - d. To begin sharing according to the .ini file configuration:

- On each CA Standard Monitor, open the Windows Services Control Panel and restart the CA ADA Data Transfer Manager service.
- On each CA Multi-Port Monitor, from the Web interface, open the Process Status page and restart the caperformancecenter_devicemanager process.

It can take up to 25 minutes for the management console to report shared WAN-optimized Client segment data.

Update the Shared Configuration

Update the shared management console configuration to add or remove a shared management console.

Follow these steps:

1. Edit the DTMDistributedConsoles.ini file.
2. Update the list of IP addresses to include each management console you want, separating each IP address on a new line. Specify each IP address in dot-decimal notation.
3. Copy the updated DTMDistributedConsoles.ini file to all monitoring devices. For a:
 - CA Standard Monitor, copy the file to the <ADA_HOME>\bin folder
 - CA Multi-Port Monitor, copy the file to the /opt/NetQoS/bin folder

If you removed a management console from the configuration file, make sure you remove the configuration file from the associated monitoring devices.

4. Begin sharing according to the updated .ini file configuration:
 - On each CA Standard Monitor, open the Windows Services Control Panel and restart the CA ADA Data Transfer Manager service.
 - On each CA Multi-Port Monitor, from the Web interface, open the Process Status page and restart the caperformancecenter_devicemanager process.

It can take up to 25 minutes to begin sharing WAN-optimized Client segment data between each management console.

Delete a Monitoring Device

When deleting a CA Standard Monitor or CA Multi-Port Monitor that shares WAE performance data, be sure to reset the configuration on the other shared monitoring devices as soon as possible so that they only share data between themselves.

Deleting a shared monitor does not automatically reset the configuration on the other shared monitoring devices. The management console continues to share data between monitoring devices and data that is shared with the deleted monitor is lost.

Follow these steps:

1. [Unassign the Cisco WAE](#) (see page 351) devices that are assigned to the CA Standard Monitor or CA Multi-Port Monitor that you want to delete.
2. [Delete the CA Standard Monitor](#) (see page 283).
3. Restart the Data Transfer Manager service on all the remaining monitoring devices. It can take up to 25 minutes to update the sharing configuration.
4. Update the Cisco WAE devices to share performance data with the available monitoring devices by restarting flow monitoring on each Cisco WAE.

To restart flow monitoring using the CLI:

- a. On the Cisco WAE, change to configuration mode by running the following command:
`config`
- b. The command line prompt changes to:
`WAE<config>#`
- c. Disable flow monitoring by running the following command:
`no flow monitor tcpstat-v1 enable`
- d. Enable flow monitoring by running the following command:
`flow monitor tcpstat-v1 enable`
- e. Return to privilege mode by running the following command:
`exit`
- f. The command line prompt changes to:
`WAE#`
- g. Check the flow monitoring status by running the following command:
`show statistics flow monitor tcpstat-v1`

Troubleshooting Tips

If the management console does not display shared data properly, validate the following:

- Within a source set, the same server IP must not be managed by more than one management console. In this case, some of the application data may appear in more than one management console, or all the data may unexpectedly appear in the management console that does not communicate with the Cisco WAE that is closest to the application server.
- The flow filters on the Cisco WAE devices are up-to-date. After you configure a monitor to share data, it can take up to 25 minutes for the monitor to communicate with each shared management console and update the flow filters on the Cisco WAE devices. The list of servers in the flow filters should be the same on all Cisco WAE devices.
- All the monitoring devices that receive packet digests from a Cisco WAE device can communicate with each other on TCP-7878.
- All the monitoring devices that receive packet digests from a Cisco WAE device can communicate with each management console on TCP-3308.
- The same configuration file is on each monitor.

Chapter 17: Monitoring with Cisco NAM

This section contains the following topics:

[How a Cisco NAM Works as a Monitoring Device](#) (see page 371)

[Add a Cisco NAM Monitoring Device](#) (see page 374)

[Edit a Cisco NAM Monitoring Device](#) (see page 378)

[Edit the NAM Monitor Feed](#) (see page 379)

[NAM Incidents](#) (see page 380)

[Delete a Cisco NAM Monitoring Device](#) (see page 381)

[Troubleshoot a Cisco NAM Monitoring Device](#) (see page 383)

How a Cisco NAM Works as a Monitoring Device

A Cisco NAM blade or appliance (Cisco NAM) serves as a type of monitoring device for CA Application Delivery Analysis and lets you troubleshoot at the IP address level with up to 30-second resolution. Using a Cisco NAM as a monitoring device reduces the server footprint for CA Application Delivery Analysis. In addition, a Cisco NAM:

- Monitors networks from switches or routers
- Reports on traffic usage
- Captures and decodes packets
- Tracks response times to pinpoint application problems to the network or server


How a Cisco NAM Works

The Cisco NAM Metric Agent creates metric digest files that include computed response time metrics for all traffic that you mirror to the Cisco NAM.

The NAM monitor feed on the CA Application Delivery Analysis Manager receives metric digest files from Cisco NAMs on its Management NIC and processes the response time statistics to automatically monitor the busiest applications and any user-defined applications.

A Distributed CA Application Delivery Analysis Manager is required.

As shown in the following diagram, you can collect packets from a SPAN, mirror port, network tap, and Cisco WAE devices. The data is processed through the Cisco NAM Metric Agent, and the computed response time statistics are sent to the CA Application Delivery Analysis Manager.



The CA Application Delivery Analysis Manager receives packet digest files on its NAM monitor feed.

How Monitor Feed Assignment Works

If a Cisco NAM is the best monitoring point, the management console automatically assigns the NAM monitor feed to the server.

More information:

[How Monitor Feed Assignment Works](#) (see page 241)

Monitoring Device Considerations

When using a Cisco NAM as a monitor, consider the following:

- The Cisco NAM Metric Agent computes response time statistics for all of the application, server, and network combinations that it observes. The CA Application Delivery Analysis Manager discards application, server, and network combinations from the Cisco NAM that do not match the server subnet, client network, and application definitions on the CA Application Delivery Analysis Manager.
- It is possible to SPAN more data to the Cisco NAM Metric Agent than it can process. SPAN carefully to minimize CPU resource consumption and enable the Cisco NAM Metric Agent to make as many accurate measurements as possible without dropping the packets it cannot process.
- You cannot use domains to separate data collection between Cisco NAM monitoring devices. All Cisco NAM monitoring devices belong to the same NAM monitor feed on the CA Application Delivery Analysis Manager.
- If you configure the Cisco NAM to use a Cisco WAE device as a data source, the Cisco NAM Metric Agent can compute response time statistics from TCP headers sent by the Cisco WAE device.
- A Cisco NAM does not temporarily store packet digest files. If a Cisco NAM cannot communicate with its assigned management console, the management console reports will contain missing data.
- The metric digest files from the Cisco NAM do not contain session-level information, therefore, the Active Sessions report is not available for a NAM monitor feed. Instead, use the Metric Receiver Counter window to see summary statistics about what the Cisco NAM sends to the CA Application Delivery Analysis Manager.
- When an application talks to a client by responding to client requests within a specified port range on the client, a Cisco NAM cannot monitor the application. The Cisco Metric Agent identifies the server that hosts an application based on the SYN-ACK response from the server port rather than the port on the client to which the server is sending TCP packets. When defining an application in the management console, if you want to monitor the application with a Cisco NAM, the application's port side must be set to Application listens on these ports.
- The management console cannot launch a packet capture investigation from a Cisco NAM, but the Cisco NAM supports triggered packet capture investigations.
- URL reporting is supported in the Cisco NAM Traffic Analyzer console rather than the management console when the Cisco NAM acts as the monitor.

More information:

[Manage User-Defined Applications](#) (see page 116)

[Troubleshoot a Cisco NAM Monitoring Device](#) (see page 383)

Add a Cisco NAM Monitoring Device

The steps required to configure a Cisco NAM monitoring device are summarized below. See the following sections for more information.

1. [Configure a Cisco NAM to export computed response time statistics to the management console](#) (see page 375).
2. Configure the Cisco NAM to export computed response time statistics to the CA Application Delivery Analysis Manager.
3. [Verify the Cisco NAM](#) (see page 376) is connected to the CA Application Delivery Analysis Manager and is receiving response time statistics from the Cisco NAM Metric Agent.
4. [Enable the NAM monitor feed](#) (see page 377) on the management console.
5. *(Optional)* To optimize the available resources on the Cisco NAM, manually configure the SPAN data source on the Cisco NAM to monitor the same networks, servers, and applications that are monitored by the management console. The Cisco NAM Metric Agent automatically computes response time metrics for all application traffic in the SPAN data source.
6. *(Optional)* [Edit monitoring device incident threshold](#) (see page 255) for inactivity.

Prerequisites

Prerequisites for adding an Cisco NAM monitoring device are listed below:

- Cisco NAM software version 4.2.x to 5.1(2).
The management console can support a Cisco Branch Router Series NME-NAM if it only processes data for the servers and applications which are local to that branch. Do not export branch NAM data to the management console unless you are certain that the only data spanned to that branch NAM is for servers that are local to that branch.
- The [monitoring device ratio](#) (see page 249) is sized properly.
- The Cisco NAM can communicate with the management console on TCP-9996.

Configure a Cisco NAM to Export Response Time Data

Configure the Cisco NAM Metric Agent to export response time data from a SPAN, Remote SPAN (RSPAN), network tap, or Cisco WAE to the CA Application Delivery Analysis Manager. Unlike a Cisco WAE, a Cisco NAM computes and sends response time metrics directly to the CA Application Delivery Analysis Manager, and the Cisco NAM computes response time data for all of the network traffic it receives.

When working with SPAN data, it is best to configure the SPAN from the host switch. We recommend that you use SPAN data so that tier-to-tier traffic can be monitored. RSPAN data can experience delays before processing within the Cisco NAM, but some cases call for its use.

After you configure the Cisco NAM Metric Agent to export computed response time statistics to the CA Application Delivery Analysis Manager, the management console automatically adds the Cisco NAM to the list of available monitoring devices.

Use the Cisco NAM Traffic Analyzer console to:

- Enable the export of computed response time metrics.
- Validate and save the active SPAN sessions you want in the running-configuration to the startup-configuration (for switches running Cisco IOS software only).
- View the SPAN data in the Cisco NAM.

Follow these steps:

1. Enable response time data export on the Cisco NAM:
 - a. Open the NAM Traffic Analyzer console.
 - b. Click the Admin tab.
The Admin tab opens.
 - c. Click System.
System Overview opens.
 - d. Click Response Time Export.
External Response Time Reporting Console Export opens.
 - e. Select Enable Export and click Apply.
 - f. Verify that the Cisco NAM is updating its configuration to export data to the management console by clicking Diagnostics on the Admin tab.
The System Alerts page displays the Cisco NAM logging messages.
2. Validate the active SPAN sessions in the running-configuration and save them to the startup-configuration:
 - a. In the NAM Traffic Analyzer user interface, click Setup, and click NAM Data Sources.

- b. Select a SPAN session from the list and click Save to save the active SPAN session in the running-configuration to the startup-configuration (for switches running Cisco IOS software only).
 - c. Repeat the previous step for each SPAN session that you want to see in the management console.
 3. View the SPAN data:
 - a. In the NAM Traffic Analyzer user interface, click Monitor, and click Server, Response Time. Sort to find the servers with the highest number of clients.
 - b. To view detailed information, select a server and click Details.

Verify the Cisco NAM is Connected to the management console

Use the management console to verify a Cisco NAM is exporting response time data to the management console.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
3. Scroll to the NAM Device List to see the last time each Cisco NAM contacted the management console.

Enable the NAM Monitor Feed

By default, the NAM monitor feed is disabled.


The management console creates a single NAM monitor feed to process metric digest files from the Cisco NAM devices in your environment. To process the metric digest files, enable the NAM monitor feed.

If you no longer intend to use a Cisco NAM as a monitoring device for the management console, you can optimize available system resources by disabling the NAM monitor feed. To disable the NAM monitor feed, [delete the CA Standard Monitor](#) (see page 283) that resides on the management console.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
3. Scroll to the ADA Monitoring Device List and browse the list to find the CA Standard Monitor with the same Management IP address as the management console.

Use the Address column to identify the Management IP address of the monitor. The CA Standard Monitor and the console share the same Management IP address.

4. Click  to edit the CA Standard Monitor with the same management IP address as the management console. If necessary, click Add ADA Monitor to add a CA Standard Monitor with the same management IP address as the management console.

Standard Monitor Properties opens.


5. Select an option to enable or disable NAM monitoring and click OK:
 - Enable NAM Monitor
 - Disable Packet Monitor

The NAM Device List displays any Cisco NAM blades or appliances that are exporting response time data to the management console.

Edit a Cisco NAM Monitoring Device

Edit a Cisco NAM monitoring device to change its monitoring device incident response and view additional details.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
3. Scroll to the NAM Device List and click  to edit a Cisco NAM.
NAM Properties opens.
4. Click Incident Response to choose a monitoring device incident response and click OK.


Edit the NAM Monitor Feed

The management console creates a NAM monitor feed to receive response time data from Cisco NAMs. Edit the NAM monitor feed to:


- Assign a particular domain. By default, a new monitor feed is assigned to the Default Domain. If you are not using domains to separate duplicate IP traffic, this is not applicable.
- Create a pair of monitor feeds. By default, a server is monitored by a single monitor feed.

The Cisco NAM Metric Agent calculates its own 5-minute response time metrics, therefore the management console does not display active session information for a Cisco NAM monitor feed.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
3. Scroll to the ADA Monitoring Device List and click  to edit the CA Standard Monitor with the same Management IP address as the management console. Note that all Cisco NAM devices are automatically assigned to the NAM monitor feed on the management console.

Standard Monitor Properties opens.

4. Scroll to Monitor Feeds and click  to edit the NAM monitor feed.
5. Assign a secondary feed or domain and click Update.

For information about setting server properties, click Help.

6. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

More information:

[Managing Tenants](#) (see page 95)

[Create a Pair of Monitor Feeds](#) (see page 243)

NAM Incidents

The management console automatically creates a Major monitoring device incident if the Cisco NAM stops sending data to the management console for more than 15 minutes. [Edit the monitoring device incident thresholds](#) (see page 255) to specify the thresholds you want.

Delete a Cisco NAM Monitoring Device


Delete a Cisco NAM monitoring device to remove it as a source of response time data. When you delete a Cisco NAM monitoring device, any servers that were pinned to the NAM monitor feed are unpinned, and another monitor feed is automatically assigned. It can take up to 10 minutes to update the monitor feed assignment.

If you have pinned servers to a NAM monitor feed, and you want to continue monitoring server traffic while the monitoring device is temporarily offline, consider the following options:

- Pin the servers to another monitor feed before you delete the monitoring device. When you move the monitoring device back online, pin the appropriate servers to the NAM monitor feed.
- Delete the Cisco NAM monitoring device. Another monitor feed is automatically assigned, but it can take up to 10 minutes to update the monitor feed assignment.

If you delete a Cisco NAM monitoring device, and the Cisco NAM is configured to export response time data, the Cisco NAM will continue to be available as a monitoring device.

Follow these steps:

1. Disable Cisco NAM Metric Agent data export on the Cisco NAM:
 - a. Open the NAM Traffic Analyzer console.
 - b. Click the Admin tab.
The Admin tab opens.
 - c. Click System.
System Overview opens.
 - d. Click Response Time Export.
External Response Time Reporting Console Export opens.
 - e. Deselect Enable Export and click Apply.
 - f. Verify that the Cisco NAM is updating its configuration to disable data export to the management console by clicking Diagnostics on the Admin tab.
The System Alerts page displays the Cisco NAM logging messages.
2. In the management console, delete the Cisco NAM from the NAM Device List:
 - a. Click the Administration page.
 - b. Click Data Monitoring, Monitoring Devices in the Show Me menu.
 - c. Scroll to the NAM Device List and click  to delete a Cisco NAM.
 - d. In the Delete Monitoring Device Confirmation, click Continue with Delete to delete the NAM monitoring device.
The Cisco NAM is removed from the NAM Device list.

Troubleshoot a Cisco NAM Monitoring Device

View the NAM counter to display information about the netflows it receives.

To view the NAM counter window, you must [log into](#) (see page 287) the management console.

Important: Before you begin, synchronize the monitoring devices. The counter windows do not display until after you have synchronized data monitoring.

Troubleshoot a Cisco NAM to identify the cause of missing report data that should have been collected by the CA Standard Monitor. Keep in mind that after you add a Cisco NAM to the management console, it can take up to 10 minutes for the management console to report application performance data from the Cisco NAM.

Unlike other monitoring devices, a Cisco NAM does not generate session-level statistics, therefore, the management console does not report them. To see summary statistics about what the Cisco NAM sends to the management console, view the NAM counter statistics on the management console.

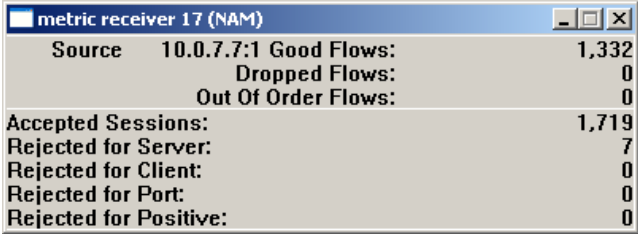
Follow these steps:

1. Log onto the management console computer or use Microsoft Remote Desktop Connection (RDC) client to remotely connect.

When using Remote Desktop to connect to a Windows Server 2003-based server, use the /admin switch to connect to the physical console session. The physical console session lets you view the feed receiver counters. For information about the /admin switch, see Microsoft KB 947723.

2. Depending on the operating system, the steps you need to take to view the NAM monitor feed statistics vary:
 - If the management console is running on Windows Server 2003, the feed receiver counters are automatically displayed. If they are not displayed, make sure you are connected to the physical console session.
 - If the management console is running on Windows Server 2008, double-click the ADA Monitor Activity shortcut on the desktop to display the feed receiver counters.

The NAM counter displays the statistics from all the Cisco NAM devices:



metric receiver 17 (NAM)	
Source	10.0.7.7:1
Good Flows:	1,332
Dropped Flows:	0
Out Of Order Flows:	0
Accepted Sessions:	1,719
Rejected for Server:	7
Rejected for Client:	0
Rejected for Port:	0
Rejected for Positive:	0

If the counter descriptions are not displayed properly, close the counter windows and reopen them by double-clicking the ADA Monitor Activity shortcut on the desktop.

If none of the feed receiver counters are displayed, verify the CA ADA Monitor service is running and verify the monitor is synchronized with the management console.

3. Interpret the NAM counter statistics to identify a problem:

Good Flows

Identifies the number of Netflow packets that were received in the order they were sent.

Dropped Flows

Identifies the number of Netflow packets that were not processed by the CA ADA Monitor service. Response time data from any metric digests that were included in the dropped Netflow packets is not included in the management console reporting.

Out of Order Flows

Identifies the number of Netflow packets that were received by the monitor, but were not received in the order they were sent.

Accepted Sessions

Identifies the number of TCP sessions that match a valid application/server/network combination on the management console.

Rejected for Server

Identifies the server IPs that did not match a server subnet monitored by the management console.

Rejected for Client

Identifies the client IP that did not match the list of client networks to be monitored by the management console.

Rejected for Port

Identifies the server port matched the list of ports to be ignored by the management console.

Rejected for Positive

Reserved for future use.

More information:

[How Client Networks Work](#) (see page 29)

[How Applications Work](#) (see page 103)

[How Servers Work](#) (see page 67)

[Perform Basic Operations](#) (see page 336)

Chapter 18: Monitoring with Riverbed Steelhead

This section contains the following topics:

[Concepts](#) (see page 387)

[Add a Monitoring Device](#) (see page 392)

[Manage a Monitoring Device](#) (see page 400)

[Troubleshoot a Monitoring Device](#) (see page 403)


Concepts

This section discusses how CA Application Delivery Analysis monitors Steelhead-optimized networks.

Introduction

A CA Standard Monitor serves as a type of monitoring device for IPv4-based Steelhead-optimized traffic. The CA Standard Monitor passively monitors WAN-optimized traffic and helps keep a continuous record of end-to-end system performance.

In the Steelhead Physical In-Path configuration below, the monitoring device at the data center monitors optimized WAN traffic across the WAN network segment.



The CA Standard Monitor in the data center can monitor both optimized traffic on the WAN segment and unoptimized traffic on the Server segment.

A separate Monitor NIC is required for optimized and unoptimized traffic.

Architecture

At the data center, the CA Standard Monitor uses separate monitor feeds to monitor both optimized and unoptimized traffic:

Steelhead monitor feed

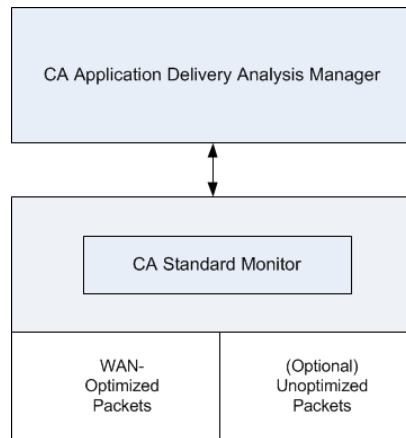
Receives WAN-optimized packets from the data center and computes response time metrics for the WAN segment of each observed application, server, and network combination.

Packet Mirror monitor feed

Receives unoptimized packets from the server switch and computes response time metrics for the Server segment of each observed application, server, and network combination. CA Application Delivery Analysis must monitor the unoptimized server traffic to report the Server network segment.

A separate Monitor NIC is required to monitor Steelhead-optimized WAN traffic and the unoptimized server traffic.

The management console automatically reports on new application traffic that matches the client networks and server subnets that are defined on the management console.



To report the response time of WAN-optimized applications at a branch location, a CA Standard Monitor is required at that branch location. The Riverbed Client Segment monitor feed receives unoptimized traffic between the client computers and the branch Steelhead appliance, and computes response time metrics for the Client segment. If a monitoring device is not deployed to a branch location, CA Application Delivery Analysis can still provide insight into WAN performance.

Network Segments

The management console generates a separate set of metrics for each of the three network segments and treats each network segment as a separate application. The management console generates a separate set of metrics for the:

- **Client segment**, which is the network segment between the clients in the branch location and the branch Steelhead appliance. Deploy a CA Standard Monitor at the branch to monitor the unoptimized Client network segment.
- **WAN segment**, which is the network segment between the branch Steelhead appliance and the data center Steelhead appliance. Deploy a CA Standard Monitor at the data center to monitor the Steelhead-optimized WAN network segment.
- **Server segment**, which is the network segment between the data center Steelhead appliance and the data center servers. Deploy a CA Standard Monitor at the data center to monitor the unoptimized Server network segment.

If necessary, use the CA Standard Monitor in the data center to monitor both the Steelhead-optimized WAN network segment and the unoptimized Server network segment.

The management console appends the network segment to the application name, for example, SMTP [Client], SMTP [WAN], and SMTP [Server]. In the example below, the SMTP application represents the unoptimized application performance. The unoptimized application response time is faster because the data comes from local users in the data center.

Performance by Application				
Application	Port(s)	Transaction Time		Observations
Wtd. Average: 7.12 ms Average: 69.29 ms				
FTP [WAN]	21	85.17 ms		6,729,198
FTP [Server]	21	82.20 ms		4,746,399
FTP [Client]	21	68.76 ms		4,765,733
FTP	21	59.94 ms		1,206
SMTP [Server]	25	39.78 ms		981,705
SMTP [Client]	25	28.26 ms		995,452
SMTP [WAN]	25	26.07 ms		1,747,521
SMTP	25	9.10 ms		10,251,044

Monitor Feed Assignment

If a monitor feed on the CA Standard Monitor is the best source for monitoring unoptimized TCP traffic on the server, the management console automatically assigns the monitor feed to the server.

Incident Thresholds for Network Segments

The management console sets different performance thresholds for each segment of a WAN-optimized network. To make the management console more or less sensitive to variations in performance, [edit the performance thresholds on each network segment](#) (see page 156).

When monitoring Steelhead-optimized traffic, the incident thresholds you set for the Client network segment apply to the branch locations where a monitoring device is deployed. A monitoring device must observe the traffic between the client computers and branch Steelhead appliance to rate application performance and create incidents.

Packet Capture Investigations

The monitoring device that monitors the unoptimized server traffic also performs packet capture investigations.

If you are using a CA Standard Monitor to monitor the unoptimized server traffic, the monitoring device:

- Launches a packet capture investigation after the management console creates an incident.
- Runs one packet capture investigation at a time.
- Captures the unoptimized Server segment. If another monitoring device is closer to the server SPAN, the corresponding monitor feed is assigned to the server.
- Copies the packet capture file from the monitoring device to the user's local computer. Depending on the size of the packet capture file, it can take a long time to open the packet capture investigation.
- The CA Standard Monitor does not provide long term packet storage.

Monitoring when Optimization Stops

Depending on how long optimization is interrupted, the management console uses different reporting methods for the affected traffic. If necessary, reset the management console to disregard recent changes in optimization and report as optimized.

Temporary Interruptions

If the monitor feed temporarily stops receiving Steelhead-optimized packets, the management console uses data from the unoptimized server SPAN to report on the application in the Server segment. Several criteria are used to determine whether an interruption is temporary or permanent, but a temporary interruption is typically less than 20 minutes.

While the management console temporarily reports the unoptimized application data from the mirrored switch port in the Server segment:

- All metrics on the Optimization page are accurate, although not all metrics populate.
- Measurements for the unoptimized sessions do not affect the Client or WAN segments.
- The Server segment is accurate because the management console uses the more accurate data from the mirrored switch port that is closest to the server.
- Network Round Trip Time [Server] on the Engineering page shows an increase because it no longer gets 100 percent local ACKs. The pass-through measurements show actual Network Round Trip Time out to the client.
- Retransmission Delay [Server] and Packet Loss Percentage [Server] can increase.

When optimization data resumes, the management console automatically reports the segmented application data in the Client, WAN, and Server segments.

Temporary interruptions to optimized monitoring occur, for example, when:

- Optimization stops. This type of interruption occurs when a Steelhead appliance does not have the resources to optimize more sessions.
- Monitoring stops. This type of interruption occurs when there is a disruption in the link between the Steelhead appliance and the monitoring device. If a monitoring device stops receiving packets on its Monitor NIC, monitoring stops.

Permanent Changes

When a monitoring device stops receiving optimized packet data for more than 20 minutes, the management console typically considers the interruption to be permanent. In this case, the management console stops reporting the unoptimized traffic in the Server segment of the application, for example, HTTP [Server]. Instead, the management console reports on the server SPAN data in the unoptimized application, HTTP. If optimization resumes, the management console automatically resumes monitoring the optimized application by network segment.

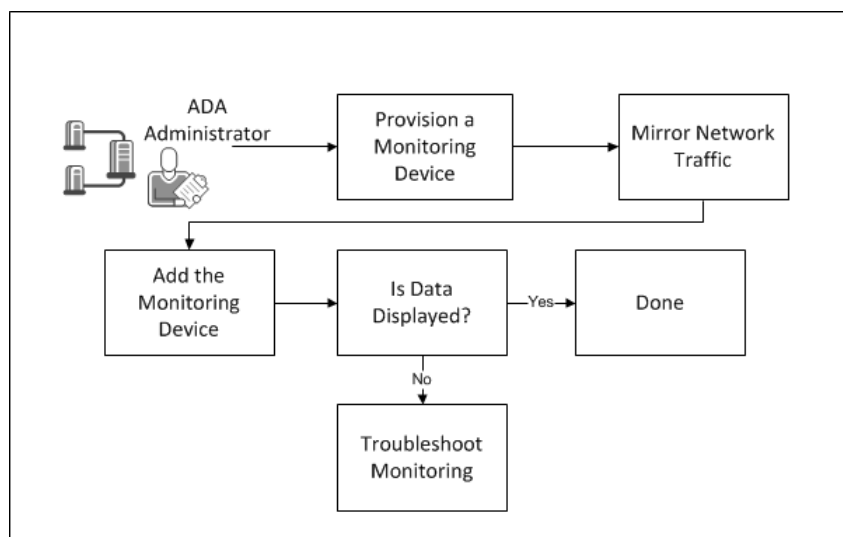
When setting up WAN optimization, for example, to measure the benefits of optimizing different applications, you do not need to wait. To report changes in optimization immediately, you can [reset optimized applications](#) (see page 347, see page 356) on the management console.

More information:

[Reset Optimized Applications](#) (see page 356)

Add a Monitoring Device

Add a monitoring device to provide visibility into WAN optimization between the data center and branch office Steelhead appliances. Monitoring WAN optimization between Steelhead appliances shows how WAN optimization affects individual application response times at each segment of the network.



Follow these steps:

1. [Provision a Monitoring Device](#) (see page 394).
2. [Mirror Network Traffic](#) (see page 394).
3. [Add the Monitoring Device](#) (see page 399).
4. (Optional) [Troubleshoot Monitoring](#) (see page 403).

Monitoring Device Considerations

When monitoring Steelhead-optimized traffic, keep in mind:

- The CA Standard Monitor monitors the following Steelhead configurations:
 - Physical In-Path
 - Virtual In-Path with WCCP (either GRE or Layer-2 redirection) configuration
- The CA Standard Monitor automatically supports the following Steelhead WAN visibility modes:
 - *Correct addressing*, uses Steelhead appliance addresses and ports over the WAN.
 - *Full Transparency*, preserves the client and server IP addresses and port numbers in the packet header fields.
- To report the WAN segment, a CA Standard Monitor in the data center is required to monitor Steelhead-optimized traffic on the WAN network segment. When configured with two NICs, up to two Steelhead appliances can be monitored with a single CA Standard Monitor.
- To report the Server segment, a monitoring device must monitor the unoptimized server traffic in the data center. If necessary, use the CA Standard Monitor in the data center to monitor the unoptimized server traffic. A separate Monitor NIC is required to monitor Steelhead-optimized WAN traffic and the unoptimized Server segment traffic.
- To report the Client segment, a CA Standard Monitor in the branch location is required to monitor traffic between the branch client computers and the branch Steelhead.
- CA Application Delivery Analysis does not monitor Web applications by URL but can monitor all HTTP traffic, for example, on Port 80.

Provision a Monitoring Device

Provision a CA Standard Monitor to monitor Steelhead-optimized traffic at the data center, and optionally, at a branch location.

Allocate memory that is based on the location of the monitoring device:

- Data Center (*recommended*): 2 GB (2048 MB)
- Branch Office (*recommended*): 1 GB (1024 MB)

When monitoring Steelhead-optimized traffic, the CA Standard Monitor in the data center can also monitor the unoptimized traffic on the data center LAN. Separate Monitor NICs are required to monitor both WAN-optimized and unoptimized traffic with the same monitoring device.

When provisioning a monitoring device, keep in mind:

- Use the CA Application Delivery Analysis setup program to install the CA Standard Monitor. Download the CA Application Delivery Analysis .iso from the CA Support Web site at support.ca.com.
- If the management console uses the Network Time Protocol (NTP), configure NTP on the monitoring device.
- The monitoring device can communicate with its assigned management console on TCP-7878.

Note: For information about provisioning a CA Standard Monitor, see the *Installation Guide* for CA Application Delivery Analysis.

Mirror Network Traffic

Use a network tap or mirror the network traffic to the CA Standard Monitor. Depending on which network segment you want, the packet mirror requirements vary.


WAN Network Segment - Physical In-Path

To report the WAN network segment, the CA Standard Monitor in the data center must receive the optimized TCP traffic between the Riverbed Steelhead appliances.

CA Application Delivery Analysis supports a Steelhead Physical In-Path configuration. Use a network tap or mirror traffic from the data center Steelhead appliance to the Monitor NIC on the monitoring device.

If CA Application Delivery Analysis does not already monitor the unoptimized Server network segment in the data center, use the CA Standard Monitor to monitor both WAN-optimized and unoptimized server traffic.

In the Steelhead Physical In-Path configuration below, the CA Standard Monitor receives a mirrored copy of the optimized traffic on the WAN network segment.



More information:

[Server Network Segment](#) (see page 398)


WAN Network Segment - Virtual In-Path

To report the WAN network segment, the CA Standard Monitor in the data center must receive the optimized TCP traffic between the Riverbed Steelhead appliances.

CA Application Delivery Analysis supports a Steelhead Virtual In-Path configuration with WCCP (either GRE or Layer-2 redirection). Use a network tap or mirror traffic from the data center Steelhead appliance to the Monitor NIC on the monitoring device.

If CA Application Delivery Analysis does not already monitor the unoptimized server traffic in the data center, use the CA Standard Monitor to monitor both WAN-optimized and unoptimized server traffic.

In the Steelhead Virtual In-Path configuration below, the CA Standard Monitor receives a mirror copy of the optimized traffic on the WAN network segment.



More information:

[Server Network Segment](#) (see page 398)

Client Network Segment

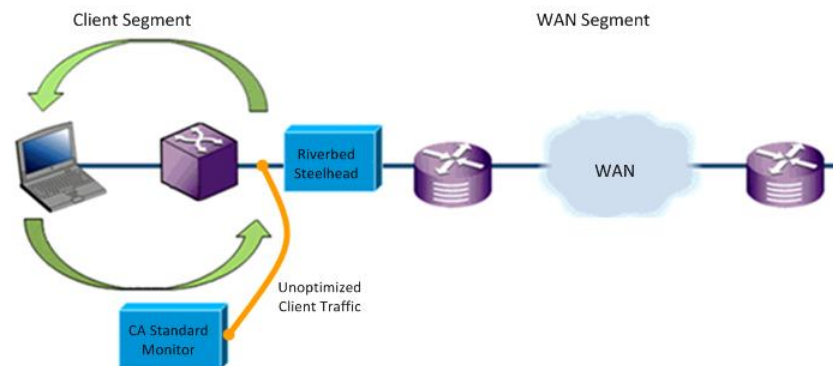
To report the Client network segment, the CA Standard Monitor in the branch office must receive the unoptimized TCP traffic between the client computers and the branch Riverbed Steelhead appliance.

Use a network tap or mirror traffic from the branch switch to the Monitor NIC on the monitoring device.

We recommend that you deploy a branch monitoring device to the remote locations where you require visibility into the list of optimized applications at a branch location. If you do not monitor the Client segment at a branch location, the CA Application Delivery Analysis user can still view a list of all applications on the WAN network segment. From this list, she can drill into the Performance Detail reports for the WAN and Server network segments.

Alternatively, in CA PC or CA NPC, create a group that contains any application with "WAN" or "Server" in it, and then filter an Engineering page report on that group and a particular client network.

Mirror the unoptimized traffic between the client computers and the branch Steelhead appliance to the monitor NIC on the monitoring device. In the Steelhead Physical In-Path configuration below, client traffic between the client computers and the branch Steelhead appliance is mirrored to the CA Standard Monitor.



Server Network Segment


To report the Server network segment, the CA Standard Monitor in the data center must receive the unoptimized TCP traffic between the data center servers and the data center server switch.

Use a network tap or mirror traffic from the server switch to a separate Monitor NIC on the monitoring device.

CA Application Delivery Analysis requires a CA Standard Monitor in the data center to monitor unoptimized traffic on the Server network segment.

If CA Application Delivery Analysis does not already monitor the unoptimized server traffic in the data center, use the CA Standard Monitor to monitor both WAN-optimized and unoptimized server traffic.

Mirror unoptimized packets between the servers and the data center switch to a separate Monitor NIC on the monitoring device. In the Steelhead Physical In-Path configuration below, the CA Standard Monitor monitors both the optimized traffic in the WAN segment and the unoptimized traffic in the Server segment.



Add the Monitoring Device

Add the CA Standard Monitor to the management console to begin reporting.

When adding a monitoring device that is not currently available on the network, after the monitor is available on the network, [synchronize](#) (see page 242) the monitoring device to establish communication between the monitor and the management console.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
3. Click Add ADA Monitor under the Show Me menu.

Standard Monitor Properties opens.

4. Complete the fields in Standard Monitor Properties. For information about specifying the monitoring device properties, click Help.

When specifying a Monitor NIC, make sure to specify the packet source. For example, if the packets are sourced from a Steelhead Physical In-Path configuration, choose the Riverbed WAN Physical In-Path option.

5. Click OK.

The monitor appears in the ADA Monitoring Device List.

6. Synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console by clicking the link.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

After 5-10 minutes, the Optimization page is displayed.

7. Click Help on the Optimization page for information about reporting on WAN-optimized applications.

If you do not see data for your application, troubleshoot the monitoring device.

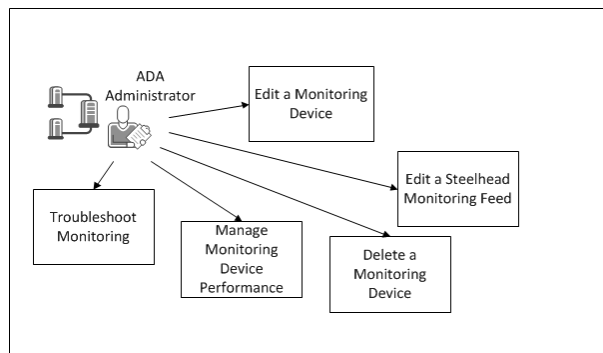
8. If you are using domains to separate duplicate IP traffic, edit the Steelhead monitor feed to assign a domain.

More information:

[Edit a Monitor Feed](#) (see page 401)

Manage a Monitoring Device

Manage a monitoring device by performing the following tasks:



Tasks

[Edit a Monitoring Device](#) (see page 400)

[Edit a Steelhead Monitor Feed](#) (see page 401)

[Manage Monitoring Device Performance](#) (see page 402)


[Delete a Monitoring Device](#) (see page 400)

[Troubleshoot Monitoring](#) (see page 403)

Edit a Monitoring Device

Edit a monitoring device, for example, to assign an incident response. While you are editing a monitoring device, you can also view any monitoring device incidents.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
3. Scroll to the ADA Monitoring Device List and click  to edit.
4. Click Properties in the third Show Me menu.

Standard Monitor Properties opens.

To view incidents for the monitoring device, click Incidents in the third Show Me menu.

5. Click Incident Response and click OK. This step selects an incident response.

More information:

[View Monitoring Device Incidents](#) (see page 254)

Edit a Monitor Feed

To assign a particular domain to a monitoring device, edit the monitor feed. By default, a new monitor feed is assigned to the Default Domain.

The management console names a monitor feed based upon its packet source:

IP address Physical Riverbed or IP address Virtual Riverbed

Indicates that the monitor feed receives Steelhead-optimized packets from the WAN network segment. For example, a monitor NIC with IP 1.1.6.44 and a Steelhead Physical In-Path configuration is named 1.1.6.44 Physical Riverbed.



IP address Packets

Indicates that the monitor feed receives unoptimized packets from the Server network segment. For example, a monitor NIC with IP 1.1.5.43 is named 1.1.5.43 Packets.

IP address Client Segment

Indicates that the monitor feed receives unoptimized packets from the Client network segment. For example, a monitor NIC with IP 1.1.6.44 is named 1.1.6.44 Client Segment.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
Scroll to the ADA Monitoring Device List, find the Standard monitor with the Packets or Steelhead monitor feed you want, and click .
3. Scroll to the Monitor Feeds section and click  to edit the monitor feed. For more information, click Help.
4. Click Update.
5. Click the link to synchronize monitoring devices with the current client network, server subnet, and application definitions on the management console.

Monitoring devices temporarily stop monitoring application performance during synchronization. To minimize interruptions to monitoring, complete all of your changes before synchronizing monitoring devices.

More information:

[Managing Tenants](#) (see page 95)

[Create a Pair of Monitor Feeds](#) (see page 243)

Manage Monitor Performance

If a monitoring device stops sending data to the management console for more than 15 minutes, the management console automatically creates a Major monitoring device incident.

More information:

[Edit Monitoring Device Incident Thresholds](#) (see page 255)


Delete a Monitoring Device

Delete a CA Standard monitoring device to remove it as a source of response time data. When you delete a monitoring device, any servers that were pinned to the corresponding monitor feed are unpinned, and another monitor feed is automatically assigned. It can take up to 10 minutes to update the monitor feed assignment.

If you have pinned servers to a monitor feed, and you want to continue monitoring server traffic while the monitoring device is temporarily offline, consider the following options:

- Pin the servers to another monitor feed before you delete the monitoring device. When you move the monitoring device back online, pin the appropriate servers to the monitor feed.
- Delete the monitoring device. Another monitor feed is automatically assigned, but it can take up to 10 minutes to update the monitor feed assignment.

Follow these steps:

1. Click the Administration page.
2. Click Data Monitoring, Monitoring Devices in the Show Me menu.
3. Scroll to the ADA Monitoring Device List and click  to delete a monitoring device.
4. Click Continue with Delete at the prompt. This step deletes the monitoring device.
The monitoring device is removed from the ADA Monitoring Device List.

Troubleshoot a Monitoring Device

Troubleshoot a monitoring device to resolve missing report data for a particular network segment.

When adding a monitoring device, [synchronize](#) (see page 242) the monitoring device to establish communication between the monitor and the management console. After you add a monitoring device, it can take up to 10 minutes to report optimized application performance on a network segment.

If the management console does not display segmented application data, troubleshoot each segment:

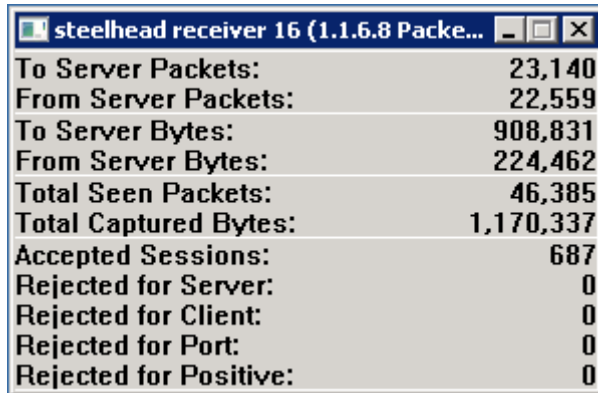
- WAN network segment:
 - View the Steelhead Receiver statistics to verify that the monitoring device has active sessions on the Steelhead monitor feed.
 - Open the Standard Monitor Properties to verify that the Monitor NIC that receives Steelhead-optimized traffic has the correct In-Path IP address of the Steelhead appliance in the data center.
 - Take a packet capture from the monitoring device Monitor port and verify that it contains the MAC address of the Steelhead appliance.
- Server network segment.
 - View the SPAN Receiver statistics to verify that the monitoring device in the data center has active sessions on the Packet Mirror monitor feed.
 - Open the Standard Monitor Properties to verify that the Monitor NIC that receives unoptimized server traffic is a Packet Mirror data source.
- Client network segment.
 - View the SPAN Receiver statistics to verify that the monitoring device in the branch location has active sessions on the Packet Mirror monitor feed.
 - Open the Standard Monitor Properties to verify that the Monitor NIC that receives unoptimized client traffic is a Client Segment Monitor data source.

View Steelhead Receiver Statistics

View Packets counter statistics for information about Steelhead-optimized TCP packet data that the CA Standard Monitor receives on a particular Monitor NIC.

Important: Before you begin, synchronize the monitoring device. The monitoring device must be synchronized to display its counter windows.

The Packets counter displays the following information:



To Server Packets:	23,140
From Server Packets:	22,559
To Server Bytes:	908,831
From Server Bytes:	224,462
Total Seen Packets:	46,385
Total Captured Bytes:	1,170,337
Accepted Sessions:	687
Rejected for Server:	0
Rejected for Client:	0
Rejected for Port:	0
Rejected for Positive:	0

To Server Packets

Identifies the total number of packets that are sent from a client to a server.

From Server Packets

Identifies the total number of packets that are sent from a server to a client.

To Server Bytes

Identifies the total number of bytes sent from a client to a server.

From Server Bytes

Identifies the total number of bytes sent from a server to a client.

Total Seen Packets

Identifies the total number of packets that match a specified application port, client network, and server subnet.

Note: If the monitor is performing normally, the number of Total Seen Packets matches the number of Received Packets. If the monitor cannot inspect all the packets that it receives, the number of Total Seen Packets falls below the number of Received Packets. In this case, the number of Dropped Packets also increases. For more information, see Dropped Packets.

Total Captured Bytes

Identifies the total byte count for packets that match a specified application port, client network, and server subnet.

Note: The CA Standard Monitor inspects each packet header to determine whether the packet matches a specified application port, client network, and server subnet. For more information, see Total Seen Packets.

Accepted Sessions

Identifies the number of TCP sessions that match a valid application/server/network combination on the management console.

Rejected for Server

Identifies the number TCP sessions where the server IP did not match a server subnet.

Rejected for Client

Identifies the number TCP sessions where the client IP did not match a client network.

Rejected for Port

Identifies the number TCP sessions where the server port matched the list of ports that the management console ignores.

Rejected for Positive

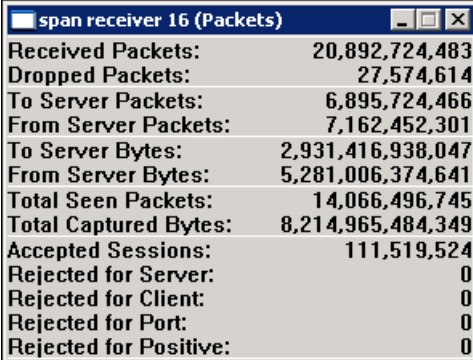
Reserved for future use.

View SPAN Receiver Statistics

View Packets counter statistics for information about unoptimized TCP packet data that the CA Standard Monitor receives on a particular Monitor NIC.

Important: Before you begin, synchronize the monitoring device. The monitoring device must be synchronized to display its counter windows.

The Packets counter displays the following information:



span receiver 16 (Packets)	
Received Packets:	20,892,724,483
Dropped Packets:	27,574,614
To Server Packets:	6,895,724,466
From Server Packets:	7,162,452,301
To Server Bytes:	2,931,416,938,047
From Server Bytes:	5,281,006,374,641
Total Seen Packets:	14,066,496,745
Total Captured Bytes:	8,214,965,484,349
Accepted Sessions:	111,519,524
Rejected for Server:	0
Rejected for Client:	0
Rejected for Port:	0
Rejected for Positive:	0

Received Packets

Identifies the total number of packets that are received, but not inspected, by the Monitor NIC on the CA Standard Monitor.

Note: Inspection of the packet header allows you to use these packets to calculate application response time metrics. For more information, see Total Seen Packets.

Dropped Packets

Identifies the total number of packets that arrived at the Monitor NIC but the packet header was never inspected. If the CA Standard Monitor is too busy processing other packets and the packet capture driver buffer is full, the packet is dropped.

To Server Packets

Identifies the total number of packets that are sent from a client to a server.

From Server Packets

Identifies the total number of packets that are sent from a server to a client.

To Server Bytes

Identifies the total number of bytes sent from a client to a server.

From Server Bytes

Identifies the total number of bytes sent from a server to a client.

Total Seen Packets

Identifies the total number of packets that match a specified application port, client network, and server subnet.

Note: If the monitor is performing normally, the number of Total Seen Packets matches the number of Received Packets. If the monitor cannot inspect all the packets that it receives, the number of Total Seen Packets is less than the number of Received Packets, and the number of Dropped Packets increases. For more information, see Dropped Packets.

Total Captured Bytes

Identifies the total byte count for packets that match a specified application port, client network, and server subnet.

Note: The CA Standard Monitor inspects each packet header to determine whether the packet matches a specified application port, client network, and server subnet. For more information, see Total Seen Packets.

Accepted Sessions

Identifies the number of TCP sessions that match a valid application, server, and network combination on the management console.

Rejected for Server

Identifies the number TCP sessions where the server IP did not match a server subnet.

Rejected for Client

Identifies the number TCP sessions where the client IP did not match a client network.

Rejected for Port

Identifies the number TCP sessions where the server port matched the list of ports that the management console ignores.

Rejected for Positive

Reserved for future use.

View Active Sessions

To view the number of active IPv4-based TCP sessions reported during the last 5-minute reporting interval, use the Active Sessions page.

Active sessions information confirms that CA Application Delivery Analysis is monitoring TCP sessions. The management console reports the number of active TCP sessions on a server by application port and network segment. For example, Port 9088 [WAN] displays the sessions information for Port 9088 application traffic on the WAN network segment.

More information:

[View Active Sessions on a Monitor Feed](#) (see page 245)

Glossary

5-minute summary file

A *5-minute summary file*, created by a CA Application Delivery Analysis Standard Monitor, Multi-Port Monitor, Virtual Systems Monitor, CA GigaStor, or a Cisco NAM, consists of 5-minute averages for each performance metric, application, server, and network [combination](#) (see page 413).

access layer

In a typical three-tier LAN network (access, distribution, core), the *access layer* is the closest layer to the server and connects servers to the network. Switches and hubs usually fall into the access layer. Typically, all server traffic can be seen at this layer, but this requires the most monitoring points.

ACK packet

During the TCP connection setup, an *ACK packet* is sent by the client to the server to acknowledge receipt of a [SYN-ACK packet](#) (see page 423) from the server.

action

See [responsive action](#) (see page 421).

application

An *application* specifies a TCP port or port range to monitor across a range of server IP addresses, such as TCP-80 traffic across a /29 server subnet.

Application Connection Time investigation (term)

An *application connection time investigation* is an [application incident response](#) (see page 411) that provides IT staff with information about how long it takes to connect to a TCP/IP application port. This includes time for the server to respond with a connection acknowledgment. A CA ADA administrator can also launch or schedule this investigation.

application incident

An application incident occurs when a Network Incident or a Server Incident impacts the performance of an application.

The threshold for a combined metric is crossed when an underlying Network Incident or Server Incident causes a combined metric for an application to cross a performance threshold.

When a combined metric crosses a threshold, the management console rates the performance impact to the application as Major (orange) or Minor (yellow), but does not create an application incident response. You must define the application incident response to launch when an underlying Network or Server Incident occurs for the application.

application incident response

An *application incident response* is an application response to a [Network incident](#) (see page 417) or a [Server incident](#) (see page 422). For example, if you configure an application incident response for the Exchange application, the management console launches the incident response when a Network incident is created by clients accessing the Exchange application, or a Server incident is created by a server that hosts the application. The management console does not launch an application incident response when the threshold for a [Combined metric](#) (see page 413), such as Data Transfer Time, is crossed. The management console lets you assign the following responses to a application: [Email notification](#) (see page 415), [SNMP Trap notification](#) (see page 423), and [Application Connection Time investigation](#) (see page 410).

availability operational level agreement (availability OLA)

An *availability operational level agreement (availability OLA)* reports the percentage of time that an application is available. For example, an application must be available on a server 99% of the time over a one month period.

baseline

A *baseline* enables you to see what is historically normal performance in your network. The management console automatically reports baselines for all TCP sessions between an application port on a server, and a client network. Use baselines to compare the current performance of an application to a historical average of past performance. A crossed baseline does not necessarily indicate a problem. Baselines are calculated hourly, and take into account hour of the day, day of the week and day of the month.

CA ADA Availability Poller service

The *CA ADA Availability Poller service* checks the availability of an application. If the server that hosts an application is monitored by a CA Standard Monitor, the check is performed by the monitoring device. Otherwise, the CA ADA Availability Poller service on the CA ADA Manager checks the availability of the application.

CA ADA Batch service

The *CA ADA Batch service* stages .dat data files for processing by the CA ADA Master Batch service on the CA ADA Manager. This service runs on the CA Standard Monitor.

CA ADA Data Pump service

The *CA ADA Data Pump service* performs weekly database maintenance on the CA ADA Manager.

CA ADA Data Transfer Manager service

The *CA ADA Data Transfer Manager service* synchronizes Cisco WAE device monitoring based on the applications, servers, and client networks defined on the CA ADA Manager. This service runs on the CA ADA Manager.

CA ADA Inspector Agent service

The *CA ADA Inspector Agent service* launches an investigation on an application, server, and its related networks. If the server that hosts an application is monitored by a CA Standard Monitor, the investigation is launched from the monitoring device. Otherwise, the CA ADA Inspector Agent service on the CA ADA Manager launches the investigation.

CA ADA Inspector service

The *CA ADA Inspector service* loads five minute .dat files processed by the CA ADA Master Batch service into the CA ADA Manager database and communicates with the CA ADA Inspector Agent service to launch investigations. This service runs on the CA ADA Manager.

CA ADA Master Batch service

The *CA ADA Master Batch service* runs on the management console to receive data files from the CA ADA Batch service on the CA Standard Monitor for processing into 5-minute .dat files. This service runs on the CA ADA Manager.

CA ADA Messenger service

The *CA ADA Messenger service* synchronizes monitoring on any assigned CA Standard Monitor, CA Multi-Port Monitor, and CA GigaStor monitoring devices, based on the applications, servers, and client networks defined on the CA ADA Manager. This service runs on the CA ADA Manager.

CA ADA Monitor Management service

The *CA ADA Monitor Management service* responds to requests from the CA ADA Manager to transfer .dat files. This service runs on the CA Standard Monitor.

CA ADA Monitor service

The *CA ADA Monitor service* receives mirrored TCP packets and packet digest files from a CA ADA monitoring device. This service runs on the CA Standard Monitor and the CA ADA Manager.

CA ADA Reader service

The *CA ADA Reader service* runs on a CA GigaStor to send packet digest files, which consist of TCP headers, to the assigned CA ADA Standard Monitor or Multi-Port Monitor for metric calculation.

CA Application Delivery Analysis Manager (CA ADA Manager)

The *CA Application Delivery Analysis Manager (CA ADA Manager)* is a component of the CA ADA architecture that provides central configuration, analysis, management, and reporting across multiple monitoring devices. The CA ADA Manager receives response time metrics from any assigned monitoring devices, including a CA ADA Standard Monitor, Multi-Port Monitor, Virtual Systems Monitor, CA GigaStor, or a Cisco NAM.

CA Observer Expert

CA Observer Expert is bundled with CA GigaStor. It combines application response time monitoring from CA ADA, with the ability to drill down into packet level data for root cause analysis.

combination

A *combination* identifies the time frame, application port, server, network, and performance metric where CA ADA computes response time metrics. For example, the management console can report on the average Network Connection Time for all applications and servers that communicated with the Development client network in the last 24 hours.

Combined metric

A *Combined metric* indicates that an application performance problem is caused by either a server that hosts the application or a network that is communicating with the application, or both. The CA ADA management console sets a performance threshold for each of the following Combined metrics: [Data Transfer Time](#) (see page 414) and [Transaction Time](#) (see page 424). Note that the management console does not create application incidents. However, because Combined metrics include both Network and Server metrics, the management console can rate a server or network as Minor (yellow) or Major (orange), and rate the corresponding performance impact on the application itself. For example, if a Server metric is rated as Minor, the management console can also rate a Combined metric for the application as Minor.

control port application

A control port application uses two TCP ports. The control port sends and receives the request information, and the data port sends and receives the actual data. The same monitoring device must monitor both the control port and the data port traffic to determine the transaction response time. Any type of monitoring device can monitor a control port application.

core layer

In a typical three tier LAN network (access, distribution, core), the *core layer* enables high speed interconnection of distribution layer devices. the core layer usually has the highest speed interconnections and the highest power routers and switches in the network. Typically only client to server transactions are seen at this layer.

Data Transfer Time

Data Transfer Time is a [Combined metric](#) (see page 413) that measures the time it takes to transmit a complete application response from the first response (the end of the [Server Response Time](#) (see page 422)) to the last packet sent in the request.

Data Transfer Time excludes the initial server response time and includes Network Round Trip Time if there is no more data to send that fits in the TCP window. The response time can be impacted by the design of the application, or the performance of the server or network.

The management console does not open an incident when the Data Transfer Time threshold is crossed.

device

A *device* can be any TCP/IP system connected to the monitored network.

discarded packet

A *discarded packet* is a packet that was intentionally discarded by a monitoring device because the packet did not match the list of applications, servers, and client networks that are specified on the management console.

distribution layer

In a typical three tier LAN network (access, distribution, core), the *distribution layer* is where routing, filtering, and policy management take place. This layer usually includes routers and layer three switches. Data is gathered when access switches send data to the distribution layer. Some server-to-server traffic can be seen at this layer, as long as the servers are on different switches.

domain

A *domain* separates client IP traffic for reporting purposes and identifies the DNS server that the management console uses to resolve the host name of a server.

dropped packet

A *dropped packet* is a packet that is not analyzed by the packet capture driver on a CA Standard Monitor or by the GigaStor Connector on a CA GigaStor, because the monitoring device is too busy to process all of the packets it receives. If the monitoring device drops too many packets, the management console creates a Major monitoring device incident. The management console does not monitor packet loss at the server switch port or the monitor NIC on the monitoring device.

Effective Network Round Trip Time

Effective Network Round Trip Time is a [Network metric](#) (see page 418) that consists of [Network Round Trip Time](#) (see page 418) plus [Retransmission Delay](#) (see page 421). Note that Retransmission Delay is not the delay due to retransmissions; it is the average amount of retransmission delay per round trip. It is important to note that the management console is adding two averages, and is actually combining two metrics.

Email notification

An *Email notification* is an [application incident response](#) (see page 411), [server incident response](#) (see page 422), or [network incident response](#) (see page 418) that notifies the recipients about a threshold violation for the affected applications, servers, or networks.

Estimated Hop Delay

Estimated Hop Delay is the estimate of how much delay was encountered between two nodes. The management console determines this estimate by using the average of all the samples taken, for example, during a [Trace Route investigation](#) (see page 424).

Expired Sessions

Expired Sessions measures the number of TCP sessions where the CA ADA Monitor service did not see the TCP session tear down (FIN or RST packet). Sessions which are inactive for a period of time are cleared out of memory and marked as Expired. The management console classifies a session as Expired if it does not observe any packets in a 15 minute period. Too many expired sessions left open can cause servers to become unresponsive.

FIN packet

In the TCP protocol, a SYN packet is used by the client when establishing a TCP connection to a server. Likewise, the *FIN packet* is used to begin the tear down or termination of the TCP connection. The monitoring device determines that a TCP conversation is being terminated when it receives a FIN or RST packet.

fragmented packet

A *fragmented packet* is a packet that has been split into multiple packets as it traverses the network.

hop

A *hop* is the logical link between two gateways in a network. When a data packet traverses a network, it will usually pass through one or more routers or gateways. The path between any two logically adjacent gateways is considered a *hop*.

incident

An *incident* is opened by the management console to raise awareness about a period of anomalous behavior on an application, server, or client network. See [responsive action](#) (see page 421).

incident response

An *incident response* helps you troubleshoot a problem at the time it occurs, and reduce the mean time to repair. Assign incident responses to your business-critical applications, servers, and networks. Incident responses notify your team about a performance degradation and actively investigate a problem to gather additional information that can help you identify the root cause of a performance degradation.

investigation

An *investigation* is an active inquiry into specific performance data on applications, networks, and servers. The management console can automatically launch an investigation in response to an incident. A CA ADA administrator can also launch or schedule an investigation.

keep-alive messages

A method to keep a TCP connection active and established persistently rather than establishing a new TCP connection for each request/response. TCP *keep-alive messages* follow a known format and do not skew response time metrics. Application keep-alives, which increment sequence numbers and contain payload, can skew measurements of some server metrics, such as Server Response Time (SRT).

Major performance rating

A *Major performance rating* is a severity state (orange) that is represented in the management console to indicate that the metric value exceeds the Major threshold. The management console sets thresholds for both Minor and Major performance degradations.

metric digest file

A *metric digest file* contains the pre-calculated response time metrics from a Cisco NAM. The CA ADA Manager receives metric digest files from a Cisco NAM.

Minor performance rating

A *Minor performance rating* is a severity state (yellow) that is represented in the management console to indicate that the metric value exceeds the Minor threshold. The management console sets thresholds for both Minor and Major performance degradations.

monitor feed

A *monitor feed* is a source of response time information, such as a CA Standard Monitor.

monitoring device

A *monitoring device* monitors TCP transactions and calculates application, server, and network response time metrics.

monitoring device incident

A *monitoring device incident* is created by the management console if the threshold for the performance and availability of a monitoring device is violated. For example, if a device becomes unreachable, no data is seen by the device, or a device discards packets.

monitoring unit

A *monitoring unit* is the processing load that is created on the CA ADA Manager by adding a monitoring device. For example, a CA Standard Monitor utilizes one monitoring unit. The CA ADA Manager supports up to 15 monitoring units.

multi-tier application

A *multi-tier application* is an application with more than one server, and communication between servers is performed by at least one server that acts as both a server to client requests, and a client to another server.

NetQoS MySql51 service

Starts and stops the MySql server which hosts the CA ADA Manager database.

Network Connection Time

Network Connection Time (NCT) is a [Network metric](#) (see page 418) that measures the amount of time between the Syn-Ack sent by the server and the Ack received back from the client. When a network is uncongested, it is a measurement of network latency that represents the minimum latency due to distance and serialization, and is the best possible round trip time for your network architecture. Sudden spikes in this value are commonly attributed to congestion, while a plateau (which goes up and stays up) typically indicates a path change.

Network incident

The management console creates a *Network incident* when the threshold for a network metric, such as Network Round Trip Time, Network Connection Time, Effective Round Trip Time, or Retransmission Delay, is exceeded during a 5-minute interval for a particular application, server, and network combination.

network incident response

An *Network incident response* is a CA ADA response to a [Network incident](#) (see page 417). The CA ADA administrator can assign an [Email notification](#) (see page 415), [SNMP Trap notification](#) (see page 423), and [Trace Route investigation](#) (see page 424) to a Network incident.

Network metric

A *Network metric* indicates an application performance problem is caused by a network that is communicating with the application. Use the CA Application Delivery Analysis management console to customize the performance thresholds for each of the following Network metrics: [Network Round Trip Time](#) (see page 418), [Network Connection Time](#) (see page 417), [Effective Network Round Trip Time](#) (see page 415), and [Retransmission Delay](#) (see page 421).

network region

A *network region* is a management console tool used to automatically expand a broad subnet definition into narrower subnets. You can define a network with up to 256 regions, such as a /16 network with 256 regions, to define 256 /24 networks. If you use network regions, the management console reports on the narrower network region subnet definitions rather than the broader network definition.

Network Round Trip Time

Network Round Trip Time is a [Network metric](#) (see page 418) that measures the time that a packet takes to travel across the network in both directions between the server and clients on a network, excluding lost packets. Application, server, and client processing time are excluded.

network tap

A *network tap* is a hardware device that lets you access the data flowing across a computer network. After a tap is in place, you can connect a monitoring device to it without impacting the monitored network. Taps enable you to view traffic occurring in both directions (upstream and downstream), but only in one broadcast domain in a switched network.

network type

A group of networks that share the same physical access to the application. For example, all of the subnets at a remote site would share the same WAN link to the data center.

observation count

The observation count measures the number of times during a five minute interval that a monitoring device calculates a performance metric for a particular application, server, and network combination. Within a TCP transaction there can be different numbers of observations for different metrics. For example, there may be more observations of Network Round Trip Time than Server Response Time. Other metrics are links and will always have the same number of observations. For example, each TCP transaction has one Server Response Time observation and one Data Transfer Time observation. To rate a metric as Normal, Minor (yellow), or Major (orange), the metric must have a minimum number of observations.

OLA

See [performance operational level agreement \(performance OLA\)](#) (see page 419) and [availability operational level agreement \(availability OLA\)](#) (see page 411).

Packet Capture investigation

A *Packet Capture investigation* is an [application incident response](#) (see page 411) or a [server incident response](#) (see page 422) that performs a filtered capture of a particular server, application port, and network experiencing a problem. A CA ADA administrator can also launch or schedule this investigation.

packet digest file

A *packet digest file* contains the TCP headers from a Cisco WAE device, or CA GigaStor Connector.

Packet Loss Percentage

Packet Loss Percentage is a [Network metric](#) (see page 418) that measures the ratio of retransmitted data to total data within the network from the vantage point of the monitoring device adjacent to the server. The monitoring device can see packets retransmitted by the server due to data loss in the server-to-client direction along the network path. When data loss occurs in the client-to-server direction before reaching the server, the monitoring device cannot observe the packet loss, and that loss is not included in the Packet Loss Percentage. On the Engineering page of the management console, Packet Loss Percentage is part of the QoS report.

performance operational level agreement (performance OLA)

A *performance operational level agreement (performance OLA)* lets you evaluate compliance with application performance goals for a remote site. By default, the management console does not define operational levels for application performance.

performance thresholds

A performance threshold is a boundary of acceptable performance behavior that exists by default for each application. Thresholds enable the management console to rate data. They contribute to incident creation, incident responses, and investigations.

Performance via SNMP investigation

A *Performance via SNMP investigation* is a [server incident response](#) (see page 422) that uses SNMP to poll a server for performance information, such as CPU and memory utilization. A CA ADA administrator can also launch or schedule this investigation.

permission set

A defined list of application, server, and network aggregations that a user has permission to view. An aggregation can be a member of one or more permission sets.

Ping Response Time investigation

A *Ping Response Time investigation* is a [server incident response](#) (see page 422) that measures the time it takes to receive a ping reply after sending a ping request, and report on the packet round trip time. A CA Application Delivery Analysis administrator can also launch or schedule this investigation.

Ping Response Time vs. Packet Size investigation

A *Ping Response Time vs. Packet Size* investigation measures the time it takes to receive a ping reply for ping requests (data packets) of various sizes. This investigation helps track excessive delays and lack of connectivity at various packet sizes. A CA ADA administrator can manually launch or schedule this investigation.

port exclusion

A *port exclusion* filters the application port traffic at the monitoring device and maximizes the available resources on the management console, while at the same time focusing the management console user's attention on the applications of interest. The monitoring device ignores TCP sessions that match a port exclusion.

product privilege

A *product privilege* determines the actions that can be performed by a user on the Administration page.

Refused Session Percentage

Refused Session Percentage is a [Server metric](#) (see page 422) that measures the percentage of connection requests the server explicitly rejects during the reporting interval. This metric is part of the Unfulfilled TCP/IP Session Requests report in the CA ADA management console.

report page

The management console organizes report data under standard *report pages* designed for particular types of users such as operations personnel, executives, and engineers.

responsive action

A *responsive action*, such as sending a notification or starting an investigation, is a response to a performance threshold violation.

Retransmission Delay

Retransmission Delay is a [Network metric](#) (see page 418) that measures the elapsed time between the original packet send and the last duplicate packet send. The management console reports Retransmission Delay as an average across observations and not just for the retransmitted packets. For example, if one packet in a set of 10 requires 300 ms of retransmission time, the Retransmission Delay is reported as 30 ms (300 ms/10 packets).

role

A *role* specifies the pages of the CA ADA management console that are displayed to a CA ADA user.

RST packet

A RST packet is a normal way to end a TCP session. A web browser typically ends a session with RST rather than FIN. The management console counts a RST packet during the connection handshake as an Unfulfilled Session Request. If the monitoring device sees a RST before the TCP three way handshake has completed, the management console considers the session to be rejected.

sensitivity level

The *sensitivity level* is a unitless measure on the scale of 0-200 that is applied to a proprietary formula which calculates a new threshold for each client, server and application combination, based on historical data. The management console automatically generates a new threshold value for a metric each night at midnight GMT, using percentile statistics from the last 30 days. The management console automatically generates a separate set of threshold values for the users who access an application from each client network.

Server Connection Time

Server Connection Time (SCT) is a [Server metric](#) (see page 422) that measures the amount of time that a server takes to acknowledge the initial client connection request by sending a Syn-Ack in response to the client's SYN packet.

Server incident

A Server incident is created by the management console when the threshold for a Server metric, such as Server Response Time, Server Connection Time, Refused Session Percentage, or Unresponsive Session Percentage, is exceeded during a five minute interval for a particular application, server, and network combination.

server incident response

A *Server incident response* is a response by the management console to a [Server incident](#) (see page 422). The management console lets you assign the following responses to a Server incident: [Email notification](#) (see page 415), [SNMP Trap notification](#) (see page 423), [Ping Response Time investigation](#) (see page 420), [Performance via SNMP investigation](#) (see page 420), and [Packet Capture investigation](#) (see page 419).

Server metric

A *Server metric* indicates that an application performance problem is caused by a server that hosts the application. Use the CA ADA management console to customize the performance thresholds for each of the following Server metrics: [Server Response Time](#) (see page 422), [Server Connection Time](#) (see page 422), [Refused Session Percentage](#) (see page 421), and [Unresponsive Session Percentage](#) (see page 424).

Server Response Time

Server Response Time is a [Server metric](#) (see page 422) that measures the time it takes for a server to send an initial response to a client request, or the initial server think time. Increases in the Server Response Time generally indicate a lack of server resources such as CPU, memory, disk I/O, a poorly written application, or a poorly performing tier in a multi-tier application.

server subnet

A server subnet identifies a contiguous range of server IP addresses that are monitored by each monitoring device. When defining an application, you can assign a particular server subnet to an application, to enable the management console to automatically monitor application performance across a contiguous range of server IP addresses.

severity

Severity classifies performance data as None, Unrated, Minor, Major, and Unavailable, by established thresholds over a time period.

SNMP profile

A *SNMP profile* is used by the management console to manage SNMPv3 user credentials and SNMPv1 and SNMPv2 community names. A SNMP profile maintains the SNMP user credentials required by the management console to query the SNMP agent on a server or network device and to send a SNMP trap message.

SNMP Trap notification

A *SNMP Trap notification* is an [application incident response](#) (see page 411), [network incident response](#) (see page 418), or [server incident response](#) (see page 422) that notifies a SNMP Manager about the Open or Closed incident status of the affected applications, servers, or networks.

SPAN

Switched Port Analyzer (SPAN), also known as port mirroring, is used on a Cisco network switch to send a copy of all network packets seen on one switch port to a network monitoring connection on another switch port. This is commonly used by network appliances to monitor network traffic. SPAN enables monitoring devices to view traffic occurring on multiple broadcast domains on one or more switch ports. SPAN capabilities vary by chassis.

SYN packet

In the TCP protocol, conversations (connections) between clients and servers are established by means of a three-way handshake. The *SYN packet* is sent by the client to a server to initiate the connection setup. A monitoring device uses the SYN packet in the timing and analysis of monitored connections on the network.

SYN-ACK packet

During the TCP connection setup, a *SYN-ACK packet* is sent by the server to the client to acknowledge receipt of a [SYN packet](#) (see page 423) from the client. A monitoring device uses the SYN-ACK packet in the timing and analysis of monitored connections on the network.

synchronize monitoring devices

Synchronize monitoring devices to monitor TCP sessions based on the current client network, server subnet, and application definitions on the management console. To minimize temporary interruptions to monitoring during synchronization, complete all changes before synchronizing monitoring devices.

tap

See [network tap](#) (see page 418).

three-way handshake

A *three-way handshake*, in the TCP protocol, is used to establish a connection between a client and a server. The [SYN packet](#) (see page 423) is sent by the client to a server to initiate the connection setup. A [SYN-ACK packet](#) (see page 423) is sent by the server to the client to acknowledge receipt of a SYN from the client. Finally, an [ACK packet](#) (see page 410) is sent by the client to the server to acknowledge receipt of a SYN-ACK from the server and to establish the TCP connection. A monitoring device uses the three-way handshake in the timing and analysis of monitored connections on the network.

thresholds

See [performance thresholds](#) (see page 419).

Trace Route investigation

A *Trace Route investigation* is a [network incident response](#) (see page 418) that records the path and each hop between a monitoring device and end-points to monitor latency and routing issues, and optionally, SNMP poll each router for its performance information. A CA ADA administrator can also launch or schedule this investigation.

traceroute

Traceroute refers to either of two types of diagnostic tools used in incident analysis: ICMP or TCP.

transaction

A *transaction* is a TCP request and all the subsequent responses. A single application transaction, such as loading a web page, can consist of multiple TCP transactions.

Transaction Time

Transaction Time is a [Combined metric](#) (see page 413) that measures the amount of time elapsed from when the client sends the request to when it receives the last packet in the response. Transaction Time is the sum of [Server Response Time](#) (see page 422), [Network Round Trip Time](#) (see page 418), [Retransmission Delay](#) (see page 421), and [Data Transfer Time](#) (see page 414). The management console does not open an incident when the Transaction Time threshold is crossed.

Unrated performance rating

An *Unrated performance rating*, indicated by a gray severity state on the Operations page of the management console, means that either there is insufficient past data to establish a threshold (two full business days of data are needed), or there were not enough observations to exceed the minimum observations threshold.

Unresponsive Session Percentage

An *Unresponsive Session Percentage* is a [Server metric](#) (see page 422) that measures the percentage of sessions where a connection request was sent, but the server did not respond. This metric is part of the Unfulfilled TCP/IP Session Requests view.

WAN

A *wide area network (WAN)* is a network that typically covers a large, diverse area comprised of multiple Local Area Networks or LANs. WANs can be private, used by different offices of a single enterprise, or public, such as the Internet.

WAN optimization device

A *WAN optimization device* reduces the volume of traffic that is transferred between the data center and a remote office through compression or other algorithms. Cisco WAE devices and Riverbed Steelhead appliances are examples of WAN optimization devices.

Index

5

5-minute summary file • 410

A

access layer • 410

ACK packet • 410

action • 410

Add a CA GigaStor Monitoring Device • 325, 327

Add a CA Standard Monitor • 272, 273

Add a CA Virtual Systems Monitor • 306

Add a Cisco NAM Monitoring Device • 374

Add a Cisco WAE Monitoring Device • 349

Add a Client Network • 44

Add a Maintenance Period to a Maintenance Schedule • 91

Add a Maintenance Schedule • 89

Add a Monitoring Device • 392

Add a Network Device • 229

Add a Network Device Group • 231

Add a Network Type • 52

Add a Port Exclusion • 109

Add a Server • 78

Add a Server Subnet • 72

Add a SNMP Profile • 227

Add an Action to a Monitoring Device Incident Response • 259

Add an Action to a Network or Server Incident Response • 177

Add an Incident Response • 175

Add an Incident Response to a Monitoring Device • 257

Add Client Networks to a Domain • 98

Add Performance Thresholds • 154

Add Servers to a Domain • 99

Add the Monitoring Device • 399

Additional SPAN Considerations • 310

application • 410

Application Connection Time investigation (term) • 410

Application Connection Time Investigations • 170

application incident • 410

application incident response • 411

Application Incident Response Trap Specifications • 187

Application Incident Responses • 167

Application Keep-Alive Messages • 136

Application Name • 183

Application Port Exclusions • 106

Architectural Overview of Application Delivery Analysis • 18

Architecture • 388

Assign a CA GigaStor to a Monitoring Device • 328

Assign a Cisco WAE to a Monitoring Device • 351

Assign a Maintenance Schedule to a Server • 93

Assign a Monitoring Device Incident Response • 261

Assign a Network Type to a Client Network • 54

Assign a Source Set to a Cisco WAE Device • 362

Assign a Source Set to a Server • 363

Assign an Incident Response • 179

Assign an Incident Response to a Network Type • 181

Assign an Incident Response to a Server • 180

Assign an Incident Response to an Application • 180

Assign Domains to Monitor Feeds • 100

Assign IP Addresses to the Network Connections • 316

Assign Servers to an Application • 126

availability operational level agreement (availability OLA) • 411

B

Back Up and Restore the Database • 223

baseline • 411

Block a CA GigaStor Input Port • 332

C

CA ADA Availability Poller service • 411

CA ADA Batch service • 411

CA ADA Data Pump service • 411

CA ADA Data Transfer Manager service • 411

CA ADA Inspector Agent service • 412

CA ADA Inspector service • 412

CA ADA Master Batch service • 412

CA ADA Messenger service • 412

CA ADA Monitor Management service • 412

CA ADA Monitor service • 412

CA ADA Reader service • 412

CA Application Delivery Analysis Manager (CA ADA Manager) • 412
CA Observer Expert • 413
CA Performance Center (CA PC) • 149
CA Technologies Product References • 3
Change the IP Address • 224
Check for Duplicate Client Networks • 297
Check Server Availability • 206
Check the management console Log File • 291
Client Network Segment • 397
ClientId • 55
ClientSetId • 55
combination • 413
Combined metric • 413
Combined Metrics • 143
Concepts • 387
Configure a Cisco NAM to Export Response Time Data • 375
Configure the Advanced Settings for Network Connections • 316
Configure the Cisco WAE to Export Response Time Data • 350
Configure the Network Connections • 314
Configure the Virtual Switch • 306
Contact CA Technologies • 3
control port application • 413
core layer • 414
Create a Control Port Application • 125
Create a CSV File • 40, 82
Create a Multi-Tier Application • 134
Create a Pair of Monitor Feeds • 243
Create a Passive Probe Instance for Each User • 326
Create a Standard Application • 117
Create a Web Application • 119, 122
Create an Application Performance OLA for a Group of Networks • 196
Create an FTP Application • 123
Create the Virtual Machine • 312

D

Data Source Synchronization • 97
Data Source System Groups • 216
Data Transfer Time • 414
Database Capacity • 247
Database Growth Control • 248
Default Client Networks • 38
Delete a CA GigaStor • 337
Delete a CA GigaStor Monitoring Device • 336

Delete a CA Standard Monitor • 283
Delete a Cisco NAM Monitoring Device • 381
Delete a Cisco WAE Monitoring Device • 355, 356
Delete a Client Network • 47
Delete a Device Group • 233
Delete a Maintenance Period • 92
Delete a Maintenance Schedule • 90
Delete a Monitoring Device • 369, 402
Delete a Monitoring Device Incident Response • 259
Delete a Network Device • 231
Delete a Network Type • 53
Delete a Port Exclusion • 112
Delete a Responsive Action • 178, 260
Delete a Scheduled Report • 234
Delete a Server • 80
Delete a Server Subnet • 75
Delete a SNMP Profile • 228
Delete a Source Set • 364
Delete a System-Defined Application • 115
Delete a User-Defined Application • 129
Delete an Application Performance OLA • 199
Delete an Incident Response • 177
DeleteNetworkDefinition • 58
Description • 55
device • 414
Disable Flow Monitoring on a Cisco WAE • 355
Disable Pop-Up Blocking • 22
Disable the Packets Monitor Feed • 284
discarded packet • 414
distribution layer • 414
domain • 414
Drill into Data Sources Permission • 213
dropped packet • 414

E

Edit a CA GigaStor Monitoring Device • 333
Edit a CA Standard Monitor • 276
Edit a Cisco NAM Monitoring Device • 378
Edit a Cisco WAE Monitoring Device • 352
Edit a Client Network • 46
Edit a Maintenance Period • 92
Edit a Monitor Feed • 401
Edit a Monitoring Device • 400
Edit a Network Device • 230
Edit a Network Device Group • 232
Edit a Network Type • 53
Edit a Port Exclusion • 111
Edit a Responsive Action • 178, 260

- Edit a Server • 79
- Edit a Server Subnet • 74
- Edit a SNMP Profile • 228
- Edit a System-Defined Application • 114
- Edit a User-Defined Application • 128
- Edit an Application Performance OLA • 198
- Edit an Incident Response • 176
- Edit Database Storage Preferences • 221
- Edit Monitoring Device Incident Response Name • 258
- Edit Monitoring Device Incident Thresholds • 255
- Edit Performance Thresholds • 150
- Edit Performance Thresholds for WAN-Optimized Network Segments • 156
- Edit Schedules for Email Reports • 233
- Edit the GigaStor Monitor Feed • 334
- Edit the NAM Monitor Feed • 379
- Edit the Packets Monitor Feed • 277
- Edit the WAN Opt Monitor Feed • 353
- Edit Thresholds for Non-Optimized Traffic on an Optimized Application • 157
- Edit Thresholds for the Optimized Client Segment • 158
- Edit Thresholds for the Optimized Server Segment • 162
- Edit Thresholds for the Optimized WAN Segment • 160
- Edit Thresholds from the Administration Page • 151
- Edit Thresholds from the Operations Page • 152
- Effective Network Round Trip Time • 415
- Eliminate Duplicate Packets on VLANs • 311
- Email notification • 415
- Email Notifications • 168
- Enable an Application Availability OLA • 208
- Enable and Disable Availability Monitoring • 256
- Enable Availability Monitoring • 204
- Enable Default Performance Thresholds for a Group of Networks • 155
- Enable the NAM Monitor Feed • 377
- Enable XFF Translation • 271
- Ensure Data Integrity and Use Anti-Virus Software • 236
- Establish Operational Levels from Historical Data • 194
- Estimated Hop Delay • 415
- Expired Sessions • 415
- Export Client Networks to a CSV File • 49
- Export Server Definitions to a CSV File • 85

F

- Filter Out Keep-Alive Messages • 281
- FIN packet • 415
- Find a Client Network • 34
- Find a Server • 77
- Find an Application • 105
- Find Unknown Client Networks with the QoS Users Report • 36
- Finish the Setup • 318
- fragmented packet • 415
- Frequently Asked Questions • 217

G

- GigaStor Incidents • 335
- Give the User Permission to a Passive Probe Instance • 331
- Group Client Networks by Network Type • 49
- Group Network Devices for Investigation • 231

H

- hop • 415
- Host Name Resolution • 71
- How /32 Client Networks Work • 70
- How a CA GigaStor Works as a Monitoring Device • 321
- How a CA Standard Monitor Works as a Monitoring Device • 265
- How a CA Virtual Systems Monitor Works as a Monitoring Device • 302
- How a Cisco NAM Works • 372
- How a Cisco NAM Works as a Monitoring Device • 371
- How an Incident Response is Launched • 166
- How Application Availability OLAs Work • 207
- How Application Availability Reporting Works • 203
- How Application Performance is Rated • 140
- How Applications Work • 103
- How Availability Monitoring Works • 201
- How Availability OLA Reporting Works • 207
- How Cisco WAAS Works • 343
- How Cisco WAAS Works as a Monitoring Device • 342
- How Client Networks Work • 29
- How Domain-Based Reporting Works • 98
- How Domains Separate Traffic • 97
- How Error Reporting Works • 61
- How Incident Responses Work • 165

How Incidents Open and Close • 147
How Maintenance Schedules Work • 88
How Monitor Feed Assignment Works • 241, 268, 323, 344, 372
How Monitor Feeds Work • 240, 268
How Monitoring Device Synchronization Works • 242
How Monitoring Devices Work • 239
How Monitoring Works when Optimization Stops • 345
How Multi-Tiered Applications Work • 131
How Network Regions Work • 32
How Network Segments Work • 344
How Network Types Work • 50
How Network-Based Reporting Works • 31
How Operational Level Metrics Work • 192
How Packet Capture Investigations Work • 269, 323
How Performance Metrics Work • 141
How Performance OLA Reporting Works • 190
How Performance OLA Thresholds Work • 191
How Performance OLAs Work • 189
How Performance Thresholds Work • 138
How Performance Thresholds Work for WAN-Optimized Network Segments • 345
How Port Exclusions Work • 107
How Priority Applications Work • 104
How Server Incidents for Application Availability Work • 203
How Servers Work • 67
How SNMP Profile Discovery Works • 226
How Source Sets Work • 361
How the CA GigaStor Connector Works • 322
How the CA Standard Monitor Works • 266
How the Management Console Manages Database Growth • 246
How the Network List Works • 30
How the Port Exclusion List Works • 108
How the Server List Works • 69
How the Server Subnet List Works • 68
How to Access the Management Console • 20
How to Configure the Cisco Nexus 1000V • 309
How to Configure the VMware vSwitch • 307
How to Maintain Hard Disk Drives • 234
How To Monitor a Multi-Tiered Application • 132
How to Navigate the Administration Page • 22
How to Set Up and Maintain CA Application Delivery Analysis • 27
How to Test the Web Services API • 60

How to Update System Security and Install Windows Updates • 235
How User Account Permissions Work • 210
How XFF Translation Works • 270

I

Import a CSV File • 42
Import Client Networks from a CSV File • 39
Import Server Definitions • 84
Import Server Definitions from a CSV File • 81
incident • 416
incident response • 416
Incident Thresholds for Network Segments • 390
Input • 184, 185, 186
InsertNetworkDefinition • 57
Install and Configure Software on the GigaStor Appliance • 326
Install CA Observer on the User's Computer • 330
Install the CA Gigastor Connector • 326
Integrated Security • 211
Internet-Facing Web Applications • 121
Interpret SNMP Traps • 187
Introduction • 387
Introduction to Tenancy • 95
investigation • 416
Issues with Domain Group Policies • 237
Issues with Third-Party Software • 236

K

keep-alive messages • 416

M

Major performance rating • 416
Manage a Monitoring Device • 400
Manage a Multi-Tiered Application • 130
Manage Client Networks • 37
Manage Client Networks with Web Service Methods • 54
Manage Console Settings • 223
Manage Incidents Using Web Service Methods • 181
Manage Monitor Performance • 402
Manage Monitoring Device Incidents • 253
Manage Monitoring Device Performance • 278
Manage Network Devices • 229
Manage Scheduled Email • 233
Manage Server Subnets • 72
Manage Servers • 76
Manage SNMP Profiles • 225

- Manage System-Defined Applications • 113
- Manage the Database • 219
- Manage User-Defined Applications • 116
- Managing Application Availability • 201
- Managing Application Performance OLAs • 189
- Managing Applications • 103
- Managing Client Networks • 29
- Managing Incident Responses • 165
- Managing Performance Thresholds • 137
- Managing Servers • 67
- Managing Tenants • 95
- Managing User Account Permissions • 209
- metric digest file • 416
- Metric Name • 183
- Minor performance rating • 416
- Mirror Network Traffic • 394
- Monitor a Server with a Group of Cisco WAE Devices • 361
- Monitor Application Availability • 205
- monitor feed • 417
- Monitor Feed Assignment • 389
- Monitored TCP Sessions • 23
- monitoring device • 417
- Monitoring Device Administration • 239
- Monitoring Device Considerations • 269, 324, 348, 373, 393
- monitoring device incident • 417
- Monitoring Device Operations • 279
- Monitoring Device Ratios • 249
- Monitoring Device Recommendations • 251
- monitoring unit • 417
- Monitoring when Optimization Stops • 390
- Monitoring with a CA GigaStor • 321
- Monitoring with a CA Standard Monitor • 265
- Monitoring with a CA Virtual Systems Monitor • 301
- Monitoring with Application Delivery Analysis • 19
- Monitoring with Cisco NAM • 371
- Monitoring with Cisco WAAS • 341
- Monitoring with Riverbed Steelhead • 387
- multi-tier application • 417

N

- NAM Incidents • 380
- Naming Conventions • 34, 77, 106
- NAT Firewall Communication • 274
- NetQoS MySQL51 service • 417
- NetQoS Performance Center (CA NPC) • 148
- Network Connection Time • 417

- Network incident • 417
- network incident response • 418
- Network Incident Responses • 166
- Network metric • 418
- Network Metrics • 142
- Network Name • 183
- network region • 418
- Network Round Trip Time • 418
- Network Segments • 389
- network tap • 418
- network type • 418
- NetworkDefinitionsList • 58
- NetworkType • 55

O

- Object Identifier Specifications • 182
- observation count • 419
- OLA • 419
- Options to Customize Performance Thresholds • 144

P

- Packet Capture investigation • 419
- Packet Capture Investigations • 171, 390
- packet digest file • 419
- Packet Loss Percentage • 419
- Parameter Descriptions • 54
- Perform Basic Operations • 252, 336
- Perform System Maintenance • 234
- Performance OLA Tips and Tricks • 193
- performance operational level agreement (performance OLA) • 419
- performance thresholds • 419
- Performance via SNMP investigation • 420
- Performance via SNMP Investigations • 172
- Permanent Changes • 347, 391
- permission set • 420
- Ping Response Time investigation • 420
- Ping Response Time Investigations • 173
- Ping Response Time vs. Packet Size investigation • 420
- Pinning a Monitor Feed to a Server • 86
- Plan for Deployment • 303
- port exclusion • 420
- Port Usage and Firewalls • 304
- Post-Installation Steps • 319
- Prerequisites • 96, 272, 325, 349, 374
- product privilege • 420
- Product Privileges • 212

Product Upgrade Support • 237
Provision a Monitoring Device • 394
Purge Data from the Database • 222

R

Recommended Browser Settings • 21
Refused Session Percentage • 421
Regions • 56
ReloadCollectors • 59
Rename a Server Maintenance Schedule • 90
Rename a Source Set • 363
Rename the Network Connections • 315
report page • 421
Request a Hardware Replacement • 237
Required Services • 220, 267
Reset Optimized Applications • 356
responsive action • 421
Retransmission Delay • 421
role • 421
Role Rights • 213
Roles • 212
RST packet • 421
Run the CA Application Delivery Analysis Setup Program • 317

S

Sample Perl Script • 62
Schedule Server Maintenance • 87
Search the Network List • 35
Secure Packet Capture Investigation Files • 275
Send Status Updates Specifications • 186
sensitivity level • 421
Sensitivity Levels (Dynamic Values) • 145
Server Connection Time • 422
Server incident • 422
server incident response • 422
Server Incident Responses • 167
Server metric • 422
Server Metrics • 143
Server Name and IP • 182
Server Network Segment • 398
Server Response Time • 422
server subnet • 422
severity • 422
Share Optimization Data between Management Consoles • 365
Share WAN-Optimized Performance Data • 367
ShowVersion • 59

Sizing Recommendations • 323, 347
SNMP Poll a Router for Client Networks • 48
SNMP profile • 423
SNMP Trap notification • 423
SNMP Trap Notifications • 169
SPAN • 423
Static Thresholds (Static Values) • 146
Strategy for Mirroring Application Conversations • 133
Subnet • 56
Support for XFF Translation • 270
SYN packet • 423
SYN-ACK packet • 423
synchronize monitoring devices • 423
Synchronize Time with a Time Server • 318
System Administration • 219
System Groups • 215
System Requirements • 305

T

Take a Packet Capture • 297
tap • 423
TCP Session Identification • 71
Temporary Interruptions • 346, 391
The Administrator Role in Application Delivery Analysis • 20
The Status of the Database • 221
three-way handshake • 423
thresholds • 424
Tips for Groups • 217
Trace Route investigation • 424
Trace Route Investigations • 174
traceroute • 424
transaction • 424
Transaction Time • 424
Troubleshoot a CA GigaStor Monitoring Device • 337
Troubleshoot a CA Standard Monitor • 285
Troubleshoot a Cisco NAM Monitoring Device • 383
Troubleshoot a Cisco WAE Monitoring Device • 357
Troubleshoot a Monitoring Device • 403
Troubleshoot Communication Issues • 263
Troubleshoot Dropped Packets • 299
Troubleshoot Incident Responses • 181
Troubleshoot Missing Data • 264, 298
Troubleshoot Monitoring • 261
Troubleshoot the CA ADA Monitor Service • 291
Troubleshooting Tips • 370

U

- Unassign a CA GigaStor • 335
- Unassign a Cisco WAE • 354
- Understanding Application Delivery Analysis • 17
- Unrated performance rating • 424
- Unresponsive Session Percentage • 424
- Update the Shared Configuration • 368
- Update Trusted Internet Sites • 22
- UpdateNetworkDefinition • 57
- Users and Groups • 214

V

- Verify Active Sessions • 286
- Verify At Least One Monitor Feed is Active • 296
- Verify Communication with the management console • 296
- Verify the Cisco NAM is Connected to the management console • 376
- Verify the Cisco WAE Configuration • 360
- Verify the NIC Settings • 294
- Verify the Registry Settings for Network Adapters • 293
- Verify the Registry Settings for the Monitor Service • 295
- Verify the Status and Priority of the NICs • 292
- View Active Sessions • 357, 407
- View Active Sessions on a Monitor Feed • 245
- View Active Sessions on the GigaStor Monitor Feed • 338
- View an Hourly Summary of Session Counts • 246
- View GigaStor Counter Statistics • 339
- View Investigations for a Network Device Group • 232
- View Monitor Feed Statistics • 287
- View Monitoring Device Incidents • 254
- View Network Device Investigations • 230
- View Sessions Information • 244
- View SPAN Receiver Statistics • 289, 406
- View Steelhead Receiver Statistics • 404
- View the List of Applications • 26
- View the List of Client Networks • 23
- View the List of Servers • 24
- View the Monitoring Device Status • 262
- View WAN Opt Counter Statistics • 358

W

- WAAS Incidents • 354

- WAN • 425
- WAN Network Segment - Physical In-Path • 395
- WAN Network Segment - Virtual In-Path • 396
- WAN optimization device • 425
- Web Application Considerations • 120
- Web Service Methods • 56
- Web Service Specifications • 184
- Welcome to Application Delivery Analysis • 17
- Why Network Types Are Useful • 51
- Why System-Defined Applications are Excluded • 202
- Windows Administrator Credentials • 219