

CA Performance Management Data Aggregator

Administrator Guide 2.4.1



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 © 2015 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 Performance Management Data Aggregator (Data Aggregator)
- CA Performance Management Data Collector (Data Collector)
- CA Performance Center

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: Product Administration 9

How to Set Up Automatic Backups of Data Repository (Single-Node and Cluster Installations)	9
Data Repository Backup Considerations	10
Configure a Data Repository Backup to a Remote Host (Single-Node and Cluster Installations)	11
Configure a Data Repository Backup to the Same Host (Single-Node and Cluster Installations)	13
Configure Data Repository	15
Restore Data Repository	18
Back Up Data Aggregator	21
Restore Data Aggregator	22
View Data Aggregator Details	23
View a List of Data Collector Installations	24
Manage Data Collector Installations	25
Rebalance the Load on Data Collector	26
Load Balancing for Data Collectors Pulling in Non-SNMP (CMM) Data	27
How to Move Data Collector to a Different Host	28
Determine the Unique Identifier for Data Collector	30
Stop Data Collector	30
Install Data Collector on a Different Host	32
Data Aggregator Configuration Changes During Network Disconnects to a Data Collector Host	34
Configure Data Collector When the Data Aggregator IP Address Changes	34
Data Collector Caching of Polled Data When the Data Aggregator Host is Unavailable	35
Calculate the Memory Required for Poll Data Caching	36
Modify the Data Cache Memory Limit	37
Data Repository Audit Process	38
Data Repository Heartbeat Monitor Process	39
Choose Another Host in a Cluster When Selected Host Fails	40
Modify Maximum Memory Usage for Data Aggregator and Data Collector Components After Installation (Optional)	41
Modify the External ActiveMQ Memory Limit After Installation (Optional)	44
Data Retention Management	45

Chapter 2: Restarting the Component Services 47

Stop and Restart Data Aggregator	47
Stop and Restart Data Collector	48
Stop and Restart Data Repository	49
Stop and Restart the ActiveMQ Broker	51

Chapter 3: Discovering Your Network 53

Device Discovery	53
Discovery Workflow	54
SNMP Profiles.....	55
Discovery and Polling	57
Discovery and Polling in VMware Environments	59
Discovery Profiles	60
View a List of Discovery Profiles.....	61
Create Discovery Profiles	61
Edit Discovery Profiles.....	65
Delete Discovery Profiles	68
Run On Demand Discoveries	68
Schedule Discoveries	70
View Discovery Results.....	71
Discovery From Other Data Sources	73
Device Type Modifications	74
Rediscovery	76

Chapter 4: Managing the Infrastructure 77

Customizing Device and Component Management Workflow	77
Monitoring Profiles	80
Factory Monitoring Profile Associations	81
View Monitoring Profiles	82
Assign or Remove Monitoring Profiles from Device Collections.....	83
caim--Configure a Monitoring Profile Poll Filter	85
Factory Device Collections	85
All Devices Device Collection	87
All Routers Device Collection	87
All Servers Device Collection	87
All Servers Device Collection	88
All Manageable Devices Device Collection	88
All ESX Hosts Device Collection	88
All Virtual Machines Device Collection	89
All VMware vCenters Device Collection	89
Custom Device Collections	89
View Monitored Devices	90
Delete Devices.....	92
Change the Primary IP Address for a Monitored Device.....	93
Delete Retired Components	94
IP Domain Deletions	96
Tenant Deletions	97

Disable Tenants	98
Enable Tenants	99
Device Reconfiguration	99
How to Manage Change Detection	100
Update Device Reconfiguration Automatically	102
Update Device Reconfiguration Manually	103

Chapter 5: Managing Interfaces 105

How to Poll Critical Interfaces Faster than Noncritical Interfaces	105
View Your Monitoring Profiles	107
Copy a Factory Monitoring Profile	108
Set an Interface Filter	110
Considerations for Interface Filters and Multiple Monitoring Profiles	111
Assign Your Monitoring Profile to a Device Collection.....	112
View Monitored Devices to Verify Results.....	113
How to Set and Activate an Interface Filter	115
Clear an Interface Filter	116
Interface Components Naming Convention	117
Interface Utilization Calculation	117
Override Speed In and Speed Out Values on Interfaces	118

Chapter 6: Eventing 119

Event Performance Guidelines.....	119
How to Monitor Event Processing	120
How to Remediate When the Threshold is Exceeded	121
Performance Management Events.....	121
Baseline Averages	122
How to Monitor Device Performance Using Events	122
Monitoring Metrics with Event Rules.....	124
Create a Custom Device Collection	126
Add Rules to a Custom Device Collection	127
Create a Monitoring Profile and Add Event Rules.....	128
Assign the Monitoring Profile to a Custom Device Collection	131
View Events	132
How to Configure Notifications from Event Manager.....	132
Event Types	135

Chapter 7: Reporting 137

How to Use Views	137
Baseline Averages	138

95th Percentile	138
Standard Deviation.....	139
Minimum and Maximum Values	139

Appendix A: Calculations **141**

Baseline Average Calculations.....	141
95th Percentile Calculations.....	146
Standard Deviation Calculations	148
Total Calculations	150
Minimum and Maximum Values	151

Appendix B: Troubleshooting **153**

Troubleshooting: Discovery Does Not Start	153
Troubleshooting: Polling Has Stopped on Discovered Metric Family	154
Troubleshooting: Polling Stopped Event Message	155
Troubleshooting: Polling Does Not Complete for My Sensitive Device	155
Troubleshooting: Unexpected Data Aggregator Shutdown	156
Troubleshooting: I am Unable to Back Up Data Repository	157
Troubleshooting: Multiple SNMP Devices Trigger Intrusions Alarms	157

Glossary **159**

Chapter 1: Product Administration

This section contains the following topics:

[How to Set Up Automatic Backups of Data Repository \(Single-Node and Cluster Installations\)](#) (see page 9)

[Restore Data Repository](#) (see page 18)

[Back Up Data Aggregator](#) (see page 21)

[Restore Data Aggregator](#) (see page 22)

[View Data Aggregator Details](#) (see page 23)

[View a List of Data Collector Installations](#) (see page 24)

[Manage Data Collector Installations](#) (see page 25)

[Rebalance the Load on Data Collector](#) (see page 26)

[Load Balancing for Data Collectors Pulling in Non-SNMP \(CMM\) Data](#) (see page 27)

[How to Move Data Collector to a Different Host](#) (see page 28)

[Data Aggregator Configuration Changes During Network Disconnects to a Data Collector Host](#) (see page 34)

[Configure Data Collector When the Data Aggregator IP Address Changes](#) (see page 34)

[Data Collector Caching of Polled Data When the Data Aggregator Host is Unavailable](#) (see page 35)

[Data Repository Audit Process](#) (see page 38)

[Data Repository Heartbeat Monitor Process](#) (see page 39)

[Choose Another Host in a Cluster When Selected Host Fails](#) (see page 40)

[Modify Maximum Memory Usage for Data Aggregator and Data Collector Components After Installation \(Optional\)](#) (see page 41)

[Modify the External ActiveMQ Memory Limit After Installation \(Optional\)](#) (see page 44)

[Data Retention Management](#) (see page 45)

How to Set Up Automatic Backups of Data Repository (Single-Node and Cluster Installations)

Situations can arise where you must back up Data Repository. For example, back up Data Repository before you upgrade Data Aggregator or before you set up automatic backups through a cron job. Backing up Data Repository gives you a copy of Data Repository to access in case there is an unexpected failure.

Important! The first time you back up Data Repository, a full backup is done. This full backup can take a considerable amount of time to complete, and depends on the amount of historical data that exists. Once an initial backup has been performed, subsequent scheduled backups will be incremental. In the case of a daily backup, an incremental backup will have to account for database activity that has occurred within the last 24 hours only (for example, amount of time that has passed since the last backup).

To perform an incremental backup after a full backup has been performed, provide the Vertica backup script with the same snapshotName and the same backup directory that you provided when you performed the full backup. If you change these names, a full backup is performed.

Vertica (the database) creates data files to store data. These files are never modified after they are created; new files are created and old ones are deleted. This approach allows you to use the standard rsync utility that supports fast file replication to another computer to back up Data Repository. For more information about rsync, see <http://everythinglinux.org/rsync/>.

To set up automatic backups of Data Repository, follow these steps:

1. [Review the backup considerations](#) (see page 10).
2. Do one of the following steps:
 - [Configure a Data Repository backup to a remote host](#) (see page 11).
 - [Configure a Data Repository backup to the same host](#) (see page 13).
3. [Configure Data Repository](#) (see page 15).

Data Repository Backup Considerations

Consider the following information before you back up Data Repository:

- You do not need to stop Data Repository or Data Aggregator when you back up Data Repository.
- Backups are stored in the location that you specify in the configuration file that you use to back up the database. The directory that contains the backup file has a subdirectory for each node that is backed up to that location. The subdirectory contains a directory with the name of the backup snapshot. The snapshot name is set using the snapshotName option in the configuration file.
- Perform incremental backups daily. We recommend performing backups during nonbusiness hours because backup processing is resource-intensive.
- You can back up Data Repository to a remote host, or you can back it up to the same host.

Note: If you back up to the same host, save the backup to a different partition than the one that is used by the catalog and data directories.

- Perform full backups weekly. The daily snapshots depend on the full backup. Restoring to any snapshot depends on the integrity of the full backup. Consider the following information about full backups:
 - Create a .ini file for each weekly full backup. The .ini file is required to restore to a particular snapshot. When a unique name is given to the .ini file, and the .ini file is run for the first time, a full backup is performed. Therefore, it is important to take note of your disk space. If the disk space is at a premium, we recommend keeping only one or two weeks of data (in addition to the current week). This solution requires an extra maintenance step of deleting the oldest week of backups as each new week begins.
 - Perform a full backup by either running the `/opt/vertica/bin/vbr.py -setupconfig` command to generate a new .ini file, or by making a copy of the current version of the .ini file. Copy the existing .ini file to a new .ini file and then change the value for “snapshotName” in the new .ini file.

More information:

[How to Set Up Automatic Backups of Data Repository \(Single-Node and Cluster Installations\)](#) (see page 9)

Configure a Data Repository Backup to a Remote Host (Single-Node and Cluster Installations)

You can back up Data Repository to a remote host.

We recommend that each Data Repository node have its own remote host for backups. For example, for a cluster environment with three Data Repository nodes, each Data Repository host requires a dedicated backup host.

Important! For cluster environments, perform the following steps on each remote host that you plan to use to back up *each* cluster node. Each node in a cluster must be backed up.

Follow these steps:

1. Open a console and log in to the computer that you plan to use as a remote backup host as the root user.
2. To create the Vertica Linux database administrator user on the remote backup host, type the following command:

```
useradd database_admin_user -s /bin/bash
```

For example:

```
useradd dradmin -s /bash/bin
```

Note: Create the same Vertica Linux database administrator user on the remote backup host that exists on the Data Repository host. Be sure that the Data Repository host and the remote backup host are not connected to LDAP or the Network Information Service (NIS) and sharing the same Vertica Linux database administrator user.

3. To set the Vertica Linux database administrator user password, type the following command:

```
passwd database_admin_user
```

For example:

```
passwd dradmin
```

4. To create Vertica directories on the remote backup host, type the following commands:

```
mkdir /opt/vertica/bin
```

```
mkdir /opt/vertica/oss
```

5. To change the owner of the Vertica directories, type the following command:

```
chown -R dradmin /opt/vertica
```

6. Log out of the remote backup host.

7. To set up passwordless ssh on the Data Repository host for the remote backup host, do the following steps:

- a. Open a console and log into the Data Repository host as the Vertica Linux database administrator user.

- b. Type the following commands:

```
ssh-keygen -N "" -t rsa -f ~/.ssh/id_rsa
```

```
cp ~/.ssh/id_rsa.pub ~/.ssh/authorized_keys2
```

```
chmod 644 ~/.ssh/authorized_keys2
```

- c. To copy the Vertica Linux database administrator user public key into the remote backup host's list of authorized keys, type the following command:

```
ssh-copy-id -i dradmin@backuphost
```

- d. Open a console and log into the remote backup host as the Vertica Linux database administrator user.

- e. To copy the Vertica rsync and python tools from the Data Repository host to the remote backup host, type the following commands:

```
scp dradmin@<drhost>:/opt/vertica/bin/rsync /opt/vertica/bin
```

```
scp -r dradmin@<drhost>:/opt/vertica/oss/python /opt/vertica/oss
```

8. Verify that the remote backup host now has the new `/opt/vertica/bin/rsync` file directory and the `/opt/vertica/oss/python` directory.
9. To create the backup directory on the remote backup host, type the following command:

```
mkdir backup_directory
```

backup_directory

Indicates the directory where you want to back up Data Repository to. Select a backup directory that is on a disk partition with a large amount of free space. If these directories are not writable by the database administrator user, give this user access to these directories through utilization of the `chown` and `chmod` commands.

Note: In a cluster installation, create the backup directories before you back up the database. You can choose a different backup directory for each host.

For example:

```
mkdir ~dradmin/backups
```

More information:

[How to Set Up Automatic Backups of Data Repository \(Single-Node and Cluster Installations\)](#) (see page 9)

Configure a Data Repository Backup to the Same Host (Single-Node and Cluster Installations)

You can back up Data Repository to the same host. In a cluster environment, you must back up each node in the cluster.

Follow these steps:

1. Log in to Data Repository as the Linux user account for the database administrator user.

Note: In a cluster installation, you can log in to Data Repository from any of the three hosts that is participating in the cluster.

2. Be sure that the Linux user account for the database administrator user is set up with a passwordless ssh key.

Note: In a cluster installation, ensure that passwordless ssh keys are set up for *each* host that is participating in the cluster.

Follow these steps:

- a. To see if a passwordless ssh key is already set up, type the following command:

```
ssh hostname ls
```

hostname

Indicates the name of the host where Data Repository is installed.

If the passwordless ssh key is set up, you are *not* prompted for a password. You do not need to do anything further.

- b. If you *are* prompted for a password, ignore the prompt and hit Ctrl+C. To set up the Linux user account for the database administrator user with a passwordless ssh key, type the following command:

```
ssh-keygen -N "" -t rsa -f ~/.ssh/id_rsa  
cp ~/.ssh/id_rsa.pub ~/.ssh/authorized_keys2  
chmod 644 ~/.ssh/authorized_keys2
```

To confirm that you are *not* prompted for a password, retype the following command:

```
ssh hostname ls
```

hostname

Indicates the name of the host where Data Repository is installed.

Important! If you do not set up the passwordless ssh key, you cannot back up Data Repository. Set up a passwordless ssh key even if you are saving the backup to the same computer.

3. To create the backup directory, type the following command:

```
mkdir backup_directory
```

backup_directory

Indicates the directory where you want to back up Data Repository to. Select a backup directory that is on a disk partition with a large amount of free space. If these directories are not writable by the database administrator user, give this user access to these directories through utilization of the chown and chmod commands.

Note: In a cluster installation, create the backup directories before you back up the database. You can choose a different backup directory for each host.

For example:

```
mkdir ~dradmin/backups
```

More information:

[How to Set Up Automatic Backups of Data Repository \(Single-Node and Cluster Installations\)](#) (see page 9)

Configure Data Repository

Configure Data Repository for automated backups.

Follow these steps:

1. Log in to Data Repository as the Linux user account for the database administrator user.

Note: In a cluster installation, you can log in to Data Repository from any of the three hosts that is participating in the cluster. However, we recommend logging in to the Data Repository host that will initiate the backups.

2. To create a reusable configuration script to use to back up and restore Data Repository, type the following command as the Linux user account for the database administrator user:

```
/opt/vertica/bin/vbr.py --setupconfig
```

Note: We recommend launching this command in the target directory for the configuration file. The Linux user account for the database administrator user must have privileges to write to that directory.

You are prompted to provide answers to various questions and statements. The list of questions and statements and a description of their typical answers are as follows:

- Snapshot name: *backup snapshot name*
- Back up vertica configurations? [y/n]: y
- Number of restore points (1): 7

Note: A restore point of 7 enables Data Repository to be restored to the most recent backup or to any of the previous 7 incremental backups. If the restore point is set to 1, you can only restore Data Repository to the most recent backup or to the previous incremental backup. The oldest backup is removed when the restore point limit is reached. To retain more than the restore point, increase the restore point or change the snapshot name in the configuration file. However, changing the snapshot name starts a new set of full backups, which can double the amount of disk space that is required for backups.

- Specify objects (no default): Do not specify a value and press Return to help ensure that all objects are backed up.
- Vertica user name (dradmin): accept default by pressing Return.
- Save password to avoid runtime prompt ? (n) [y/n]: y
- Password to save in vbr config file (no default): Enter the password when you are prompted.

Note: This password must correspond to the database password for the database administrator account within Vertica.

- Backup host name (no default): *the host name for the backup*

Note: If you are backing up a cluster, you are prompted for the hostname that corresponds to each node in the cluster. You must back up each node in a cluster.

- Backup directory (no default): *the directory path where you want to back up Data Repository to*

Note: If you are backing up a cluster, you are prompted for a backup directory for each node in the cluster. You must back up each node in a cluster.

- Config file name (snapshot name.ini): accept default by pressing Return.

Verify that you have write permissions to the directory where you are creating the .ini file. If you do not enter a full path to the .ini file, the file is saved to the directory where you ran the `/opt/vertica/bin/vbr.py --setupconfig` command.

Important: The configuration file that is generated contains a clear text password.

- Change advanced settings? (n) [y/n]:n

A message indicates that the vbr configuration has been saved to a configuration file named snapshot name.ini.

3. Back up Data Repository. Type the following command:

```
/opt/vertica/bin/vbr.py --task backup --config-file  
configuration_directory_path_filename
```

configuration_directory_path_filename

Indicates the directory path and filename of the configuration file you created previously. This file is located where you ran the backup utility (`/opt/vertica/bin/vbr.py`).

For example:

```
/opt/vertica/bin/vbr.py --task backup --config-file  
/home/vertica/vert-db-production.ini
```

If you are prompted about the authenticity of the host, answer yes.

Note: In a cluster installation, you only have to perform this step on one of the hosts that are participating in the cluster.

Data Repository is backed up.

4. (Optional) If you do not want to retain the Data Repository password in clear text for future manual backups, do the following steps:

- a. Verify that the following line exists under the [Database] section:

```
dbPromptForPassword = True
```


- b. Remove the following line from the [Database] section:

```
dbPassword = password
```

Note: For automated backups, the dbPassword line must remain in the configuration file with a corresponding password. Set the dbPromptForPassword to False.

5. Do the following to set up an automated daily backup (recommended) of Data Repository:

- a. Open your preferred text editor to create a new wrapper shell script.
- b. The contents of the wrapper shell script should contain the following single line:

```
/opt/vertica/bin/vbr.py --task backup --config-file  
configuration_directory_path_filename
```

configuration_directory_path_filename

Indicates the directory path and filename of the configuration file you created previously. This file is located where you ran the backup utility (/opt/vertica/bin/vbr.py).

For example:

```
/opt/vertica/bin/vbr.py --task backup --config-file  
/home/vertica/vert-db-production.ini
```

- c. Save the contents to a new file named backup_script.sh in a location of your choice.

For example:

```
/home/vertica/backup_script.sh
```

- d. Change permissions for running the script by typing the following command:

```
chmod 777 location_backup_script.sh/backup_script.sh
```

For example:

```
chmod 777 /home/vertica/backup_script.sh
```

- e. As the Linux user account for the database administrator user type the following command:

```
crontab -e
```

- f. Add a cron job that will run the backup script that you created previously.

Note: We suggest that you create a cron job to run the script daily at an off-peak time.

For example:

```
00 02 * * * /home/vertica/backup_script.sh >/tmp/backup.log 2>&1
```

This example cron job will run the backup script every day at 2:00 AM.

Important! The first time you back up Data Repository, a full backup is done. This full backup can take a considerable amount of time to complete, and depends on the amount of historical data that exists. Once an initial backup has been performed, subsequent scheduled backups will be incremental. In the case of a daily backup, an incremental backup will have to account for database activity that has occurred within the last 24 hours only (for example, amount of time that has passed since the last backup).

Restore Data Repository

You can restore Data Repository after it was backed up. This procedure assumes that the database administrator user is part of the sudoers file.

Note: Usually, you restore Data Repository to the same computer where you backed it up from. However, you *can* restore Data Repository to a different computer. The computer that you restore to must be configured in the same way that the computer you backed up Data Repository from is. In a cluster environment, each computer you restore to must be configured in the same way that each computer you backed up each Data Repository node from is.

The following configurations have to be the same:

- the IP address
- the hostname
- the catalog and data directories
- the catalog and data directory permissions
- the Vertica Linux database administrator user credentials
- the database administrator user account credentials
- the database user account credentials

Follow these steps:

1. Stop all Data Collector hosts that are associated with Data Aggregator by logging in to the computers where Data Collector is installed as the root user or a sudo user with access to a limited set of commands. Open a command prompt and type the following command:

```
service dcmd stop
```

Data Collector hosts stop.

2. Stop Data Aggregator by logging in to the computer where Data Aggregator is installed as the root user or a sudo user with access to a limited set of commands. Open a command prompt and type the following command:

```
service dadaemon stop
```

Note: For information about creating a sudouser with access to a limited set of commands, see the *Data Aggregator Installation Guide*.

Data Aggregator stops.

3. Log in to the database server you use for Data Repository as the database administrator user, *not* as the root user.
4. Type the following command:

```
/opt/vertica/bin/adminTools
```

The Administration Tools dialog opens.

5. Select (4) Stop Database.
6. Press the Space bar next to the database name, select OK, and press Enter. You are prompted for the database password.
7. Enter the database password and press Enter.

Data Repository stops.

Note: If Data Repository does not stop, select (2) Stop Vertica on Host from the (7) Advanced Tools Menu.

8. Select Exit and press Enter.
9. To prepare to restore the Data Repository backup, log in as the Linux user account for the database administrator user to the database server you use for Data Repository.

When you set up automatic backups of Data Repository, you configured the configuration file with a restore point of seven. Data Repository can be restored to the most recent backup or to any of the previous seven incremental backups.

10. Do one of the following steps:
 - a. To restore Data Repository to the most recent backup, type the following command:

```
/opt/vertica/bin/vbr.py --task restore --config-file  
configuration_directory_path_filename
```

configuration_directory_path_filename

Indicates the filename and directory path of the configuration file you created when you ran the backup configuration procedure. This file is located where you ran the backup utility (/opt/vertica/bin/vbr.py).

For example:

```
/opt/vertica/bin/vbr.py --task restore --config-file  
/home/vertica/vert-db-production.ini
```

Note: In a cluster installation, you can run the restore task from any of the hosts that are participating in the cluster.

- b. To restore Data Repository to any of the previous seven incremental backups, type the following command:

```
/opt/vertica/bin/vbr.py --task restore --config-file  
configuration__directory_path_filename --archive_name
```

configuration__directory_path_filename

Indicates the filename and directory path of the specific configuration file you want to restore a specific archive from. You created this configuration file when you ran the backup configuration procedure. This file is located where you ran the backup utility (/opt/vertica/bin/vbr.py).

archive_name

Indicates the name of the specific restore point that you want to restore to. Change to the backup directory that the configuration file for the restore point indicates. All of the restore points that are available are listed. Determine the archive name for the restore point that you want to restore to.

For example:

```
/opt/vertica/bin/vbr.py --task restore --config-file myconfig.ini --archive  
20131020_170018
```

Note: In a cluster installation, you can run the restore task from any of the hosts that are participating in the cluster.

11. Restart Data Repository by logging in to the computer where Data Repository is installed as the database administrator user, *not* as the root user. Open a command prompt and do the following steps:
 - a. Type the following command:

```
/opt/vertica/bin/adminTools
```

The Administration Tools dialog opens.
 - b. Select (3) Start Database.
 - c. Press the Space bar next to the database name, select OK, and press Enter.
You are prompted for the database password.
 - d. Enter the database password and press Enter.
Data Repository starts.

e. Select Exit and press Enter.

12. Restart Data Aggregator by logging in to the computer where Data Aggregator is installed as the root user or sudo user with access to a limited set of commands. Type the following command:

```
/etc/init.d/dadaemon start
```

Data Aggregator starts.

13. Restart all Data Collector hosts that are associated with Data Aggregator:
- Select Admin, Data Source Settings, and click a Data Aggregator data source.
 - Click Data Collectors from the System Status menu.
 - Select all Data Collector hosts that are associated with Data Aggregator and click Start.

Data Collector hosts start.

Back Up Data Aggregator

Situations can arise where you must back up Data Aggregator. For example, back up Data Aggregator and Data Repository before you upgrade. Backing up these components gives you a copy of your settings and custom certifications to access in case there is an unexpected failure.

You do not need to stop Data Repository, Data Collector, or Data Aggregator services when you back up Data Aggregator.

Backups are stored in the location you specify, which can be on the Data Aggregator system or a different backup host system.

Note: You must have root or sudo privileges to perform this task.

Follow these steps:

- Open a command prompt.
- Use the following command to create a backup directory in a secure location on the same or different backup host system:

```
mkdir DA_Backup
```

DA_Backup

Specifies the directory path and name of the backup directory.

3. Create subdirectories within DA_Backup using all of the following commands:

```
mkdir DA_Backup/deploy_backup
mkdir DA_Backup/MIBDepot_backup
mkdir DA_Backup/CustomDeviceType_backup
```

4. Run the following commands to back up the files on the DA:

- This command backs up the custom vendor certifications. Do not back up the local-jms-broker.xml and the README files from that directory.

```
cp Data Aggregator installation
directory/apache-karaf-2.3.0/deploy/im.ca.com.*.xml
DA_Backup/deploy_backup
```

- This command backs up all the custom MIBs in the MIBDepot directory:

```
cp Data Aggregator installation directory/apache-karaf-2.3.0/MIBDepot/*
DA_Backup/MIBDepot_backup
```

- This command backs up all the custom device subtype xml files:

```
cp Data Aggregator installation
directory/apache-karaf-2.3.0/custom/devicetype/DeviceType.xml
DA_Backup/CustomDeviceType_backup/
```

Data Aggregator installation directory

Specifies the Data Aggregator install directory.

Default: /opt/IMDataAggregator

Restore Data Aggregator

You can restore the Data Aggregator information that you backed up. If Data Repository remains intact, you can restore only the Data Aggregator component.

You do not have to stop Data Aggregator before restoring. The backed up files can be dropped in the right directories even when Data Aggregator is running.

Note: You must have root or sudo privileges to perform this task.

Follow these steps:

1. Open a command prompt.
2. (Optional) In the situations where the Data Aggregator karaf service is not running, uninstall the existing Data Aggregator and reinstall it.

3. Run all of the following commands:

```
cp DA_Backup/deploy_backup/*.* Data Aggregator installation
directory/apache-karaf-2.3.0/deploy/
cp DA_Backup/MIBDepot_backup/*.* Data Aggregator installation
directory/apache-karaf-2.3.0/MIBDepot/
cp DA_Backup/CustomDeviceType_backup/*.* Data Aggregator installation
directory/apache-karaf-2.3.0/custom/devicetype/
```

If prompted, overwrite the existing file.

DA_Backup

Specifies the directory path and name of the backup directory.

Data Aggregator installation directory

Specifies the Data Aggregator install directory.

Default: /opt/IMDataAggregator

4. Wait for a few minutes for Data Aggregator to synchronize automatically with CA Performance Center. When the connections between the Data Aggregator and the Data Collector hosts are established, the Data Collector hosts resume polling.

Data Aggregator is restored.

Note: If you must restore Data Collector to a previous state, you can uninstall and reinstall Data Collector.

View Data Aggregator Details

You can view the number of manageable and pingable devices that Data Aggregator is monitoring.

The administrator can view the total number of manageable and pingable devices that Data Aggregator monitors for all tenants. Individual device totals for each tenant are also displayed in a table.

Tenant administrators can view the total number of manageable and pingable devices that Data Aggregator monitors for their tenant.

You can also view the version and the build number of Data Aggregator.

Follow these steps:

1. Open CA Performance Center as an administrator.
2. Select Admin, Data Source Settings, and click a Data Aggregator data source.
3. Click Data Aggregator from the System Status menu.

The Data Aggregator List page opens. The total number of manageable and pingable devices by tenant is displayed, and the version and the build number of the selected Data Aggregator installation is displayed.

View a List of Data Collector Installations

You can view a list of available Data Collector installations and can change some of their settings. The Data Collector List shows the tenant and IP domain to which each Data Collector installation is assigned and the Data Collector status and version. You can also see the number of devices and components that each Data Collector installation is polling, and the total number of devices that are assigned to that Data Collector instance, including devices that are not currently polled.

The administrator can see a list of Data Collector installations for all tenants. Tenant administrators can see only the Data Collector installations that are assigned to their tenant.

Follow these steps:

1. Open CA Performance Center as an administrator.
2. Select Admin, Data Source Settings, and click a Data Aggregator data source.
3. Click Data Collectors from the System Status menu.

The Data Collector List page opens, displaying a list of available Data Collector installations.

More information:

[Manage Data Collector Installations](#) (see page 25)

Manage Data Collector Installations

The administrator must select an IP domain and a tenant for each Data Collector installation. Each Data Collector instance can be associated with only one IP domain; the Data Collector instance that is associated with that IP domain carries out discovery requests.

IP domains are logical groupings that identify data from different devices and networks. Monitoring by domain means that IP addresses with associated interfaces or applications that belong to separate customer networks are monitored separately. When combined with appropriate permissions, IP domains are monitored from a single console, but users view data only for the domains that they monitor.

A *tenant* represents a customer environment that a managed service provider administers. Each tenant environment is independent and effectively functions as a separate instance of CA Performance Center. Each instance can contain multiple users and roles that are not shared among tenants.

The Default Tenant represents the tenant space for the managed service provider within the managed infrastructure. Assign the Default Tenant if you are not deploying multi-tenancy. In a single-tenant environment, the Default Tenant is the space used for monitoring the entire infrastructure.

Follow these steps:

1. Open CA Performance Center as an administrator.
2. [Navigate to the Data Collectors page](#) (see page 24).
3. Select a Data Collector instance from the list.
4. Verify that Data Collector is available for assignment. The Polled Items column lists the number of polled devices and components that are assigned to this Data Collector instance.

Important! If the number of polled devices and components is greater than one, you cannot change the tenant or IP domain assignment for the Data Collector instance.

5. Click Assign.

The Assign Data Collector dialog opens.

6. Select the tenant that you want to assign to this Data Collector instance from the drop-down list.

All monitored devices and components that this Data Collector instance discovers are automatically associated with this tenant.

Select 'Default Tenant' if you want to use the default tenant.

7. Select the IP domain that you want to associate with this Data Collector instance.

All managed devices and components that this Data Collector instance discovers are automatically associated with this IP domain.

8. Click Save.

The tenant and IP domain are assigned to the Data Collector installation.

Rebalance the Load on Data Collector

As a Data Collector instance monitors more devices, the Data Collector capacity can be exceeded and the Data Collector can become overloaded. You can transfer the work load from one overloaded Data Collector instance to other Data Collector instances. You can rebalance the load on Data Collector in two ways:

- Select the overloaded Data Collector instance and then select 'Rebalance'. The product automatically rebalances the load with another available Data Collector instances.
- Move selected devices from one Data Collector instance to another.

Important! We recommend that you do not rebalance the load on Data Collector or move a large number of items from one Data Collector instance to another during peak hours because it may impact enduser performance.

Follow these steps:

1. Open CA Performance Center as an administrator.
2. Select Admin, Data Source Settings, and click a Data Aggregator data source.
3. Click Data Collectors from the System Status menu.

You can see the number of devices and components that each Data Collector installation is polling. You can also see the total number of devices that are assigned to each Data Collector instance, including devices that are not currently polled.

Automatically Rebalance the Load on Data Collector

1. Select the Data Collector instances that you want to rebalance and click Rebalance.

Note: Be sure to select Data Collector instances within the same IP domain. Only Data Collector instances within the same IP domain can rebalance devices between one another.

2. A confirmation dialog displays the current device and polled item count for each selected Data Collector and the proposed resulting device and polled item counts.

Note: Devices can only be moved to Data Collector instances that can contact them.

3. Click Yes.

Note: Rebalancing polled items restarts the baseline average calculations for all rebalanced polled items.

Move Selected Devices to a Specific Data Collector Instance

1. Select the Data Collector instance that you want to move selected devices from.
2. In the Devices table, select the devices that you want to move to another Data Collector instance and then click Move Devices.
3. The Move Devices to Selected Data Collector dialog opens.
4. Select the Data Collector instance that you want to move your selected devices to from the drop-down list.

Note: Only Data Collector instances that are within the same IP domain are included for selection.

5. Click Yes.

Note: Moving devices restarts the baseline average calculations for the moved devices.

Load Balancing for Data Collectors Pulling in Non-SNMP (Camm) Data

Data Collector load balancing by moving devices and components from one Data Collector instance to another only applies to devices and components being monitored via SNMP or ICMP. For Data Collector instances that are pulling in Non-SNMP data via Camm that require a rebalancing of resources, you can do this by distributing Device Pack engines to other hosts in the environment. Here are the instructions on how to perform this rebalancing.

1. Install a Local Controller (LC) on new server and point to the appropriate Multi Controller (MC) server during the Installation
2. After LC is successfully installed in new server check that Camm web shows two LCs.
 - a. Open CammWEB -
 - b. Click on Hosts – Installed LC (new server) should be visible
3. Using CammWEB Select the new server and deploy the devicepack engines to be migrated
4. Log in to the MC server and navigate to:
`$Camm_INSTALL/MC/repository/<OLD_SERVER_IP>/COMPONENTS` directory

5. Execute the following:

```
'cp -R ENGINE_<devicepack>  
$CMM_INSTALL/MC/repository/<NEW_SERVER_IP>/COMPONENTS/'
```
6. If the devicepack to be migrated uses sftp/ftp/copy mechanism as data acquisition then
 - a. create following directories under
\$CMM_INSTALL/LC/repository/COMPONENTS/ENGINE_<devicepack>/ in
NEW_SERVER
 - tmp directory under
\$CMM_INSTALL/LC/repository/COMPONENTS/ENGINE_<devicepack>/tmp/input/inventory
 - input directory under
\$CMM_INSTALL/LC/repository/COMPONENTS/ENGINE_<devicepack>/tmp/input/performance
 - b. Copy the following files from OLD_SERVER to NEW_SERVER
 - \$CMM_INSTALL/COMPONENTS/ENGINE_<devicepack>/tmp/input/inventory/.historyFile.Inventory to
\$CMM_INSTALL/LC/repository/COMPONENTS/ENGINE_<devicepack>/tmp/input/inventory
 - \$CMM_INSTALL/COMPONENTS/ENGINE_<devicepack>/tmp/input/performance/.historyFile.Performance to
\$CMM_INSTALL/LC/repository/COMPONENTS/ENGINE_<devicepack>/tmp/input/performance
7. Start the devicepack from CMMWEB.

How to Move Data Collector to a Different Host

Data Collector is a component of Data Aggregator. You can move Data Collector to a different host system without having to rediscover network devices and components or lose historical data. For example, if you are a tool administrator, your server administrator can instruct you to relocate Data Collector to another host. Data Collector is polling 500,000 devices and components so you do not want to lose data or perform rediscovery.

The Data Collector component can be moved even if you have device packs installed.

Note the following considerations:

- The amount of data loss is equal to the amount of time that has elapsed from the time the old Data Collector component is shut down to the time when the new Data Collector component has been deployed.
- If the old Data Collector component happens to start accidentally, this results in the double polling of SNMP data. You see a warning similar to the following in the Data Aggregator karaf log:

```
WARN | Session Task-810 | 2013-01-02 13:52:09,062 | DCHeartBeatLog |  
ore.collector.interfaces |  
| HeartBeat message not received. Expected: 93, Received: 255
```

To fix this problem, stop or uninstall the old Data Collector component.

The following diagram shows how to move Data Collector to a different host:

Move Data Collector to a Different Host



To move Data Collector to another system, follow this process:

1. [Determine the unique identifier for Data Collector](#) (see page 30).
2. [Stop Data Collector](#) (see page 30).
3. (for CA Mediation Manager integration only) [Migrate device packs](#) (see page 32).
4. [Install Data Collector on a different host](#) (see page 32).

Determine the Unique Identifier for Data Collector

Determine the unique identifier for Data Collector before moving this component to another host.

Retrieve the Data Collector ID using *one* of the following methods:

- Log in to CA Performance Center as a user with the Administrator role, and do the following steps:
 - a. Select Admin, and then select a Data Aggregator data source from the menu.
 - b. The Data Aggregator Admin UI opens.
 - c. Select System Status, Data Collectors from the menu.
 - d. Find the Data Collector component that you want to move and notate its ID.

- Open a web browser and issue the following web service call:

`http://DA_hostname:port/rest/dcms`

DA_hostname:port

Specifies the Data Aggregator host name and the port number.

Default port: 8581

Find the <DataCollectionMgrInfo> section whose HostName and IPAddress match the one you want to move. Notate the value for <DcmID>.

Next, stop the Data Collector services on the current host.

Stop Data Collector

Stop the Data Collector services on the current host before moving Data Collector to another host.

Follow these steps:

1. If you have installed device packs for this Data Collector, take the following steps. If no device packs are installed, proceed to Step 2.

- a. Log in to CA Performance Center as a user with the Administrator role.
- b. Select Admin, and then select a Data Aggregator data source from the menu.

The Data Aggregator Admin UI opens.

- c. Select EMS Integration Profiles from the Monitoring Configuration menu.
- d. Right-click on a profile that is associated with this Data Collector host, and select Stop. Do this step for every EMS profile that is related to this Data Collector host.
- e. Archive CA Mediation Manager artifacts by running this command:

```
tar -zcvf filename
/opt/IMDataCollector/apache-karaf-{n.n.n}/MediationCenter
```

filename

Specifies the name of the archive file.

Note: This archive file is moved to the new Data Collector host later on.

2. Log on to the Data Collector host and run the following command:

```
/etc/init.d/dcmd stop
```

3. Verify Data Collector has stopped:

- a. Log in to CA Performance Center as a user with the Administrator role.
- b. Select Admin, and then select a Data Aggregator data source from the menu.
- c. Select System Status, Data Collectors from the menu.
- d. Verify the Data Collector status shows "Not Connected".

Next, you install Data Collector on the new host.

Install Data Collector on a Different Host

After you stop Data Collector services on the old host, you install Data Collector on a new host. The Data Collector data from the old host is exported to the new host during this procedure.

Follow these steps:

1. (For integration with CA Mediation Manager only) Migrate your device packs. On the old Data Collector host, run the `$CAMM_HOME/tools/migratePMtoCAMM` script with the `-t` flag.

This step assumes that you are running the script on a Data Collector server where a Local Controller is installed. You must also have the CA Mediation Manager Console running on another server.

Note: Migrated device packs are copied to `$CAMM_HOME/MigratedIMDevicepacks` in the form of .zip files. For more information about migrating device packs, see the "How to Migrate Device Packs" scenario.

2. Log in to the new host system and open a command shell session.
3. Set an environment variable with the ID you copied previously by running this command:

```
export DCM_ID=data collector id
```
4. Install Data Collector from the same session by running the **install.bin** binary.
5. Install the CA Mediation Manager LC on the same server.
6. If you have previously installed device packs for this Data Collector, perform these additional steps:
 - a. Copy the zip files that you created previously with the migration script to local directories on this host.
 - b. Use the CA Mediation Manager web Console to deploy these device packs and start them.

Note: You do *not* need to redeploy the certification packs to the Data Aggregator host.

Note: After a few polling cycles, verify that data is being collected by the new Data Collector host.

As a best practice, uninstall the old Data Collector and delete any associated EMS profiles after verifying that data is being collected on the new host. This best practice is optional.

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 is proprietary information of CA and may not be copied, transferred, reproduced, disclosed, modified or duplicated, in whole or in part, without the prior written consent of CA.

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 © 2015 CA. All rights reserved. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

Data Aggregator Configuration Changes During Network Disconnects to a Data Collector Host

Occasionally, the connection between a Data Aggregator host and a Data Collector host breaks, such as, when a network disconnect occurs. If the Data Aggregator and Data Collector processes are running during a disconnect, you can make configuration changes to the Data Aggregator installation. In this case, polling continues on the Data Collector host according to the configuration that existed before the network disconnect. Once the connection between the Data Aggregator and the Data Collector hosts reestablishes, Data Collector downloads the new configuration and adjusts polling accordingly.

For example, you make one of the following configuration changes:

- Change the expression an SNMP vendor certification uses to calculate a value on a metric family.
- Change the metric family to poll a new operational metric.

When the connection between the Data Aggregator and the Data Collector hosts is broken, the changes cannot take effect. After reconnection, Data Collector begins polling the new SNMP MIB objects used in the new expression or in calculating the new operational metric.

Configure Data Collector When the Data Aggregator IP Address Changes

To allow a Data Collector to communicate with a Data Aggregator, configure the Data Collector when you change the IP of the Data Aggregator.

Note: If the Data Collector uses the hostname, it is only necessary to restart the Data Collector to keep the communication between the Data Collector and the Data Aggregator. Configure the Data Collector only if it uses an IP address to communicate with the Data Aggregator.

Follow these steps:

1. Stop the Data Collector, if it is running. Open a command prompt and type the following command:

```
/etc/init.d/dcmd stop
```

2. Edit the hostname or address in the following file:
`/opt/IMDataCollector/apache-karaf-2.3.0/etc/com.ca.im.dm.core.collector.cfg`
Edit the following line:
`collector-manager-da-hostname`
Save the file.
3. Update the IP address in the following file:
`/opt/IMDataCollector/apache-karaf-2.3.0/jms/local-jms-broker.xml`
4. Delete the following file:
`/opt/IMDataCollector/apache-karaf-2.3.0/deploy/local-jms-broker.xml`
5. Delete the cache. Open a command prompt and type the following command:
`rm -rf /opt/IMDataCollector/apache-karaf-2.3.0/data/cache/*`
6. Start the Data Collector. Open a command prompt and type the following command:
`/etc/init.d/dcmd start`
7. Make sure that the correct address appears in the Data Collector List.
 - a. Open CA Performance Center as an administrator.
 - b. Select Admin, Data Source Settings, and click a Data Aggregator data source.
 - c. Click Data Collectors from the System Status menu.
 - d. The IP address of each Data Collector appears under the 'Address' column.
 - e. The status of each Data Collector is "Collecting Data".

Data Collector Caching of Polled Data When the Data Aggregator Host is Unavailable

Occasionally, the network connection between the Data Aggregator and the Data Collector hosts is lost. In this case, Data Collector continues polling and caches polled data in memory, up to a configurable limit. When the Data Aggregator host becomes available, the cached polled data is sent to Data Aggregator.

Polled data is processed in a "first in, first out" order. That is, the oldest cached polled data is sent to Data Aggregator first. If the cache memory limit is reached, any new polled data is lost until the Data Aggregator host becomes available and the Data Aggregator host has processed 9 percent of the cached data.

Important! Memory usage increases dramatically on the Data Collector system when Data Aggregator is unavailable.

Your memory storage requirements vary and depend on the following factors:

- The number of devices and components polled
- The polling rate
- How much data you want to retain when the Data Aggregator host is unavailable

The default value for the cache memory limit is one half of the maximum Data Collector process memory. You configured the maximum memory usage when you installed Data Collector or after the installation.

Data Collector requires a dedicated amount of memory to function properly. In a small-scale environment, with Data Collector polling 50,000 devices and components at a five-minute poll rate, 2 GB of memory is required for basic operation. In a large-scale environment, with Data Collector polling 500,000 devices and components at a five-minute poll rate, 24 GB of memory is required for basic operation. The remaining memory can be used for caching polled data.

Calculate the Memory Required for Poll Data Caching

The amount of memory that is required for caching polled data depends on the following information:

- The scale of your environment.
- How long you want to store data when the Data Aggregator host is unavailable.

To calculate the amount of memory that is required for data caching, use the following formula:

Cache Required(GB)= (Time To Cache Data(sec)× Number of Polled Items)/(262144 ×Average Poll Rate(sec))

Example: Calculate the Memory Required for Caching Polled Data for an Hour

- Calculate the memory that is required in a small-scale environment where Data Collector is polling 50,000 devices and components at a five-minute poll rate. You want to cache the polled data for an hour while Data Aggregator is unavailable:

Cache Required (GB)=(3600 ×50000)/(262144 ×300)

Cache Required (GB)=2.3 GB

Note: This calculation is in addition to the basic operational memory requirements. A small-scale environment requires 2 GB for basic operational memory requirements. Therefore, the total memory that is required is 4608 M (2 GB + 2.3 GB).

- Calculate the memory that is required in a large-scale environment where Data Collector is polling 500,000 devices and components at a five-minute poll rate. You want to cache the polled data for an hour while the Data Aggregator is unavailable:

Cache Required (GB)=(3600 ×500000)/(262144 ×300)

Cache Required (GB)=22.9 GB

Note: This calculation is in addition to the basic operational memory requirements. A large-scale environment requires 24 GB for basic operational memory requirements. Therefore, the total memory that is required is 47 GB (24 GB + 22.9 GB).

Modify the Data Cache Memory Limit

You can modify the amount of data that Data Collector caches when Data Aggregator is unavailable.

Follow these steps:

1. [Calculate the amount of memory that is required for data caching](#) (see page 36).
2. Make a note of the amount of memory that is required for data caching.
3. Log in to the computer where Data Collector is installed. Log in as the root user or a sudo user with access to a limited set of commands.

Note: For more information about the sudo user, see the *Data Aggregator Installation Guide*.

4. Stop Data Collector using this command:

```
service dcmd stop
```

5. Modify the IM_MAX_MEM memory setting for Data Collector:

- a. Access the *Data Collector installation directory*/apache-karaf-2.3.0/jms/local-jms-broker.xml file.
- b. Change the IM_MAX_MEM limit to two times the value you noted in step 2. Verify that this value is not more than the available RAM on the system.

6. Modify the cache memory limit for the JMS broker on Data Collector:
 - a. Access the *Data Collector installation directory*/apache-karaf-2.3.0/jms/local-jms-broker.xml file.
 - b. Locate the following line:

```
<memoryUsage limit="value"/>
```

value

Is the current cache limit setting.
 - c. Modify the current cache limit setting with the value that you calculated previously and save the file.
7. Make Data Collector aware of the changes to the jms/local-jms-broker.xml file. Type the following command to deploy a fake .lock file. A fake .lock file makes Data Collector think that there was an ungraceful shutdown:

```
echo `date` > /opt/IMDataCollector/apache-karaf-2.3.0/.lock
```
8. Restart Data Collector using this command:

```
service dcmd start
```

The cache memory limit is configured.

Data Repository Audit Process

The audit process audits the database daily at 3:00 AM to calculate the total space that Data Aggregator data occupies. The process estimates the size of the database using the Vertica function, 'audit'. When estimating the database, it does not include data stored in temp tables, data that has been marked for deletion but is not purged from the database, and data in Vertica monitoring tables.

CA Technologies has a license agreement with Vertica which states that the total data stored in Data Repository cannot exceed 32 TB.

To view the most recent result of an audit, access the following URL in your browser:

<http://hostname:port/rest/datarepositorymaintenance/audit>

This URL returns XML. The "Current Size" tag displays the current size of Data Repository in bytes.

Important! Review the audit results periodically. If you see a value greater than 32 TB, you are not in compliance with the license agreement. Contact CA Technical Support for further instructions.

Data Repository Heartbeat Monitor Process

The heartbeat monitor process checks whether Data Repository is up and running every 10 seconds. If the heartbeat process fails to confirm that the database is up in 5 minutes, Data Aggregator shuts down. An audit message is logged in the *Data Aggregator installation directory/apache-karaf-2.3.0/shutdown.log* file.

In a cluster environment, all nodes in the cluster are continuously checked for availability every 10 seconds. If a node cannot be contacted within 5 minutes, an event is generated and logged on the Data Aggregator device in CA Performance Center. An audit message is logged in the *Data Aggregator installation directory/apache-karaf-2.3.0/shutdown.log* file.

If the Data Repository node that failed is the primary node (through which all Data Aggregator queries were made), Data Aggregator automatically switches to the next available Data Repository node. An event is generated and logged on the Data Aggregator device.

Important! Certain administrative functions that are occurring during a high availability failover are interrupted and then fail. One poll cycle is lost. These functions will not resume after Data Repository connects to another node in the cluster environment. Administrative functions that you perform after Data Repository connects to another node in the cluster environment work as designed.

If all Data Repository nodes fail in a cluster environment, Data Aggregator is shut down.

Loss of contact with Data Repository can result in a loss of data by Data Aggregator. Resolve any connectivity or Data Repository issues before you restart Data Aggregator. Data Aggregator shuts down automatically if it fails to connect to Data Repository on start-up. To minimize data loss, the Data Collector installations continue to collect and store data locally for a time until Data Aggregator is restarted.

To recover a node that has failed, select the “Restart Vertica on Host” option on the main menu of the admintools utility and follow the prompts. Data Aggregator will not establish a heartbeat on the failed node until you restart the Vertica process on that node and there is a successful network connection.

Choose Another Host in a Cluster When Selected Host Fails

If the database host that is specified during Data Aggregator installation fails at runtime, Data Aggregator shuts down automatically. If you installed Data Repository in a cluster, point database connections to another host in the cluster before you restart Data Aggregator.

Follow these steps:

1. Open the *Data Aggregator installation directory*/apache-karaf-2.3.0/etc/dbconnection.cfg file on the Data Aggregator host.
2. Modify the following line in the dbconnection.cfg file. Modify the line to reference a hostname or IP address of one of the Data Repository cluster hosts that is still up and running:

```
dbUrl=jdbc:vertica://database server hostname:database server  
port/databasename?use35CopyFormat=true&BinaryDataTransfer=false
```

database server hostname:database server port

Indicates the hostname or IP address of Data Repository and the Data Repository port number that you entered during the Data Aggregator installation.

Default port number: 5433

Example:

If host2 is up and running in the cluster and you choose database connections to point to host2, your updated dbUrl entry could look like the following line:

```
dbUrl=jdbc:vertica://host2:5433/mydatabasename?use35CopyFormat=true&BinaryDataTransfer=false
```

3. Save the dbconnection.cfg file.
4. To restart Data Aggregator, type the following command:

```
/etc/init.d/dadaemon start
```
5. To help ensure that Data Aggregator is not still running, type the following command:

```
Ps -ef | grep java | grep -v grep
```

Data Aggregator processes are not returned when Data Aggregator is not running.

Database connections point to the specified host in the cluster going forward.

If more than one host in the Data Repository cluster fails, Data Repository and Data Aggregator shuts down automatically. The Data Repository cluster is only capable of losing one host.

If a single host in the cluster that is *not* specified during the Data Aggregator installation disconnects from the network (for example, because a firewall was put in place, or the Ethernet cable was removed), Data Aggregator shuts down. Data Aggregator restarts automatically if you set up the automatic recovery of the Data Aggregator process during the Data Aggregator installation. Once the host that is offline becomes available, return that host to the cluster. Select the “Restart Vertica on Host” option on the main menu of the admintools utility and follow the prompts.

Note: For information about setting up the automatic recovery of the Data Aggregator process, see the *Data Aggregator Installation Guide*.

If a single host in the cluster that is *not* specified during the Data Aggregator installation is stopped through the “Kill Vertica Process on Host” option on the Advanced Menu of the admintools utility, Data Aggregator continues to function. Once the host that is offline becomes available, return that host to the cluster. Select the “Restart Vertica on Host” option on the main menu of the admintools utility and follow the prompts.

Modify Maximum Memory Usage for Data Aggregator and Data Collector Components After Installation (Optional)

The default maximum memory usage for the Data Aggregator and the Data Collector components is not sufficient. To run effectively in a large-scale deployment, modify the maximum memory usage for Data Aggregator and for Data Collector. This modification can be done during or after the installation process. By default, the memory usage for Data Aggregator and Data Collector is 2 GB.

Important! The memory modifications in this procedure assume that Data Aggregator and Data Collector are installed on separate computers. This procedure also assumes that those computers are dedicated only to the installation of these components.

Follow these steps:

1. Open a console and type the following command:

```
more /proc/meminfo
```

The total memory usage is displayed.

2. Make a note of this total memory.

3. Modify the maximum memory for Data Aggregator by performing the following steps:

- a. Access the *Data Aggregator installation directory*/apache-karaf-2.3.0/bin/setenv file.
- b. Modify the IM_MAX_MEM=*number unit* line for large-scale deployments.

number unit

Indicates the maximum amount of memory. *number* is a whole, positive number, and *unit* is "G" or "M". Subtract 2 GB from the total memory you noted previously and enter it here. 2 GB are reserved for other operating system operations.

For example: 33544320 KB - 2G = 30 GB

IM_MAX_MEM=30G

For example:

IM_MAX_MEM=4G

- c. Save the file.
- d. Restart Data Aggregator using the following command:

```
service dadaemon start
```

Data Aggregator starts and synchronizes with CA Performance Center automatically.

- e. In order for the memory setting change to persist during a Data Aggregator upgrade, modify the /etc/DA.cfg file, replacing the updated value for the property "da.memory".

For example:

```
da.memory=4G
```

4. Modify the maximum memory for all Data Collector hosts by performing the following steps:

- a. Access the *Data Collector installation directory*/apache-karaf-2.3.0/bin/setenv file.
- b. Modify the IM_MAX_MEM=*number unit* line for large-scale deployments.

number unit

Indicates the maximum amount of memory. *number* is a whole, positive number, and *unit* is "G" or "M". Subtract 2 GB from the total memory you noted previously and enter it here. 2 GB are reserved for other operating system operations.

For example: 33544320 KB - 2G = 30 GB

IM_MAX_MEM=30G

For example:

IM_MAX_MEM=4G

- c. Save the file.
- d. Restart Data Collector hosts using the following command:

`service dcmd start`
- e. In order for the memory setting change to persist during a Data Collector upgrade, modify the /opt/DCM.cfg, replacing the updated value for the property "IM_MAX_MEM".

For example:

IM_MAX_MEM=4G

The maximum amount of memory is configured for large-scale deployments.

Example: Configure the Maximum Memory Usage for Data Aggregator After You Install Data Aggregator

The following example configures the maximum memory usage for Data Aggregator where the total memory is 3354432 KB:

1. Open a console and type the following command:

```
more /proc/meminfo
```

The following result displays:

```
MemTotal: 3354432KB
```

2. Calculate the maximum memory that is required for large-scale deployments:

Equation: total memory - 2G = maximum memory for large-scale deployments

Solution: 3354432 KB - 2G = 30G
3. Access the *Data Aggregator installation directory*/apache-karaf-2.3.0/bin/setenv file.
4. Modify the IM_MAX_MEM=*number unit* line for large-scale deployments:

IM_MAX_MEM=30G
5. Save the file.
6. Restart Data Aggregator.

The maximum amount of memory is modified for large-scale deployments.

Modify the External ActiveMQ Memory Limit After Installation (Optional)

The Data Aggregator installer calculates the memory that is needed on your system to accommodate the ApacheMQ process. However, you can manually modify the memory limit settings to fine tune ActiveMQ on your Data Aggregator system. For example, you can modify the settings under the following circumstances:

- When the system memory has changed.
- When the number of Data Collector systems have changed.
- To optimize the memory settings.
- When you have determined that ActiveMQs performance is degraded, by monitoring either the JConsole or the CA Performance Management custom chart with ActiveMQ metrics.

Follow these steps:

1. Calculate the amount of memory for ActiveMQ based on the following settings:

Maximum java heap size

This value is set to 20% system memory by default. The minimum value is 512M.

Initial minimum java heap size

This value should be 50% of maximum java heap size.

Memory limit for all messages

This value should be 50% of the maximum java heap size.

Memory limit per queue

This value should be calculated based on how many Data Collector installations you have.

Example: The memory per queue

$(\text{system memory for all messages})/5/(\text{Data Collector count})$

2. Log in to the computer where Data Aggregator is installed. Log in as the root user or a sudo user with access to a limited set of commands.

Note: For more information about the sudo user, see the *Data Aggregator Installation Guide*.

3. Type the following command to stop the ActiveMQ broker:

```
/etc/init.d/activemq stop
```

4. Modify the java heap size for ActiveMQ:
 - a. Access the **activemq** file under `broker/apache-activemq-version/bin`.
 - b. Locate the line that defines `ACTIVEMQ_OPTS_MEMORY`.
 - c. Change `-Xms` to be the Initial minimum java heap size.
 - d. Change `-Xmx` to be the Maximum java heap size.
 - e. Save the file.
5. Modify the ActiveMQ memory limit for the producer flow control:
 - a. Access the `activemq.xml` file in the *Data Aggregator installation directory*/`broker/apache-activemq-version/conf` file.
 - b. Locate the following line and change the value to Memory limit for all messages:

```
<memoryUsage limit="value"/>
```
 - c. Locate the following line, change the value to Memory limit per queue:

```
<policyEntry queue=">" producerFlowControl="true"
memoryLimit="value"/>
```

Note: For more information, refer to <http://activemq.apache.org/producer-flow-control.html>
<http://activemq.apache.org/producer-flow-control.html>.
6. Type the following command to start the ActiveMQ broker:

```
./etc/init.d/activemq start
```

Your new settings are activated.

Data Retention Management

Data retention rates for Data Repository are manageable. The default data retention rates in Data Repository are set to conserve the disk space and improve reporting for most users. Polled data is generated every 5 minutes by default for a given device, and this data represents the most granular data available in the product. This raw, polled data is set to roll up at an hourly interval. The rolled up data is an aggregation of polled values, which provides a higher level, less detailed view in reports. Daily and weekly rollups can be kept longer than polled or hourly data because they require less disk space to store.

However, you can change the rate at which Data Repository retains the polled data, hourly rollup data, daily rollup data, and weekly rollup data. For example, you can change the polled data retention value to 30 days to conserve disk space. Find the balance that best suits your needs and environment.

Note: For information about how to change data retention rates, see the *REST Web Services Guide*.

By default, data is retained in Data Repository for the following number of days:

- Polled data: 45 days

Note: If you upgraded to this release from a previous release of Data Aggregator, the polled data retention will not change from the prior default of ten days.

- Hourly rollup data: 90 days
- Daily rollup data: 365 days
- Weekly rollup data: 730 days

The minimum number of days that Data Repository can retain data for is as follows:

- Polled data: 2 days
- Hourly rollup data: 8 days
- Daily rollup data: 31 days
- Weekly rollup data: 366 days

Chapter 2: Restarting the Component Services

This section contains the following topics:

[Stop and Restart Data Aggregator](#) (see page 47)

[Stop and Restart Data Collector](#) (see page 48)

[Stop and Restart Data Repository](#) (see page 49)

[Stop and Restart the ActiveMQ Broker](#) (see page 51)

Stop and Restart Data Aggregator

Situations can arise in which you must stop and restart Data Aggregator. For example, the operating system of the Data Aggregator host requires an upgrade. Stop Data Aggregator, perform the actions that you want to take, and restart Data Aggregator. Data Aggregator then resumes processing.

During a planned Data Aggregator shutdown, all polled data that is received before Data Aggregator shuts down are sent to Data Repository. This polled data is preserved for reporting and other purposes.

Consider the following information about data loading, rollups, and event threshold processing when you plan to stop Data Aggregator:

- All data that has been received from Data Collector components for the current poll cycles is processed before Data Aggregator stops. The data is not lost.
- If the threshold event processing has started at the time of shutdown, the data processing that has been received from Data Collector components completes before Data Aggregator stops.
- Threshold event processing resumes when Data Aggregator restarts.
- If the rollup processing has started at the time of shutdown, rollup processing for the data that has been collected from the Data Collector hosts completes before Data Aggregator stops.
- Rollup processing resumes when Data Aggregator restarts.

Data Aggregator could shut down in an unplanned manner, such as when the computer where Data Aggregator is installed loses power. In this case, Data Aggregator stops abruptly. In this case, polled data and threshold event information can be lost. Data loading of queued data from the Data Collector hosts resumes when Data Aggregator restarts. Event threshold processing and rollup processing for queued data resumes when Data Aggregator restarts.

Follow these steps:

1. Log in to the computer where Data Aggregator is installed. Log in as the root user or a sudo user with access to a limited set of commands.

Note: For more information about the sudo user, see the *Data Aggregator Installation Guide*.

2. Open a command prompt and do one of the following steps:
 - a. If you are logged in as the root user, type the following command:

```
service dadaemon stop
```
 - b. If you are logged in as the sudo user, type the following command:

```
sudo service dadaemon stop
```

Polling continues on Data Collector if it is running and polling when Data Aggregator is stopped. Data Collector queues polled data for future delivery to Data Aggregator.

3. Relocate the computer, or perform any other administrative tasks.
4. Log in to the computer where Data Aggregator is installed. Log in as the root user or a sudo user with access to a limited set of commands.

Note: If you installed Data Aggregator as the sudo user, you set up a sudo command alias for the `/etc/init.d/dadaemon` command. Use the sudo command to run the `dadaemon start` script. For more information about the sudo user, see the *Data Aggregator Installation Guide*.

5. Open a command prompt and type the following command:

```
service dadaemon start
```

Data Aggregator starts and synchronizes with CA Performance Center automatically.

When Data Aggregator starts, any queued, polled data on the Data Collector host are sent to Data Aggregator. The newest data is discarded if the queued data exceeds a disk space limit that is configured on the Data Collector system. As a result, there is a gap in the polled reporting data.

Stop and Restart Data Collector

Situations can arise when you must stop and restart Data Collector. For example, the computer where Data Collector is installed can lose power or can lock up. Or, you want to relocate the computer. In this case, stop and restart Data Collector. If you want to install an operating system patch, stop and start Data Collector.

Follow these steps:

1. Log in to the computer where Data Collector is installed. Log in as the root user or a sudo user with access to a limited set of commands.

Note: For more information about the sudo user, see the *Data Aggregator Installation Guide*.

2. Open a command prompt and do one of the following steps:
 - a. If you are logged in as the root user, type the following command:

```
service dcmd stop
```
 - b. If you are logged in as the sudo user, type the following command:

```
sudo service dcmd stop
```

When Data Collector is stopped, all ongoing polling stops. You cannot run any discoveries.

3. Relocate the computer, or perform any other administrative tasks.
4. Start Data Collector by logging in to the computer where Data Collector is installed. Log in as the root user or a sudo user with access to a limited set of commands.

Note: If you installed Data Collector as the sudo user, you set up a sudo command alias for the `/etc/init.d/dcmd` command. Use the sudo command to run the dcmd start script. For more information about the sudo user, see the *Data Aggregator Installation Guide*.

5. Open a command prompt and type the following command:

```
service dcmd start
```

When Data Collector is restarted, scheduled polling resumes. You can resume running discoveries. Data Collector resynchronizes with CA Performance Center automatically.

Stop and Restart Data Repository

Situations can arise when you must stop and restart Data Repository. For example, the computer where Data Repository is installed can lose power or can lock up. Or, you want to relocate the computer. In this case, stop and restart Data Repository. If you want to install an operating system patch or you want to upgrade to a new version of Data Repository, stop and restart Data Repository.

Follow these steps:

1. Log in to the computer where Data Aggregator is installed. Log in as the root user or a sudo user with access to a limited set of commands.

Note: For more information about the sudo user, see the *Data Aggregator Installation Guide*.

2. Open a command prompt and type the following command:

```
service dadaemon stop
```

3. Log in to the database server you use for Data Repository as the database administrator user, *not* as the root user.

4. Type the following command:

```
/opt/vertica/bin/adminTools
```

The Administration Tools dialog opens.

5. Select (4) Stop Database.
6. Press the Space bar next to the database name, select OK, and press Enter.
You are prompted for the database password.
7. Enter the database password and press Enter.

Data Repository stops.

Note: If Data Repository does not stop, select (2) Stop Vertica on Host from the (7) Advanced Tools Menu.

8. Select Exit and press Enter.
9. Relocate the computer, or perform any other administrative tasks.
10. Log in to the database server you use for Data Repository as the database administrator user, *not* as the root user.

11. Type the following commands:

```
/opt/vertica/bin/adminTools
```

The Administration Tools dialog opens.

12. Select (3) Start Database.
13. Press the Space bar next to the database name, select **OK**, and press Enter.
You are prompted for the database password.
14. Enter the database password and press Enter.
The database starts.
15. Select (E) Exit and press Enter.

16. Start Data Aggregator by logging in to the computer where Data Aggregator is installed. Log in as the root user or a sudo user with access to a limited set of commands.

If you installed Data Aggregator as the sudo user, you set up a sudo command alias for the service dadaemon command. Use the sudo command to run the dadaemon start script.

Note: For more information about the sudo user, see the *Data Aggregator Installation Guide*.

17. Open a command prompt and typing the following command:

```
service dadaemon start
```

Data Repository restarts.

Stop and Restart the ActiveMQ Broker

Restart the Apache ActiveMQ broker if Data Aggregator detects a problem with ActiveMQ and Data Aggregator is unable to restart the broker successfully. You can also manually stop and restart the service when needed.

Follow these steps:

1. Open the following directory from the command line:

```
cd da_install_dir/broker/apache-activemq-version/bin
```

da_install_dir

Specifies the location of the Data Aggregator installation directory.

apache-activemq-version

Specifies the version of Apache ActiveMQ.

Example: apache-activemq-5.5.1b

2. Run the stop command:

```
./activemq stop -jmxurl  
service:jmx:rmi:///jndi/rmi://localhost:11099/jmxrmi --jmxuser  
admin --jmxpassword activemq da_broker
```

-jmxurl service:jmx:rmi:///jndi/rmi://localhost:11099/jmxrmi

Specifies the location of the activemq broker. This location only changes when a user modifies the port, or a user externalizes the broker to another system.

Note: Modifying the port number is supported, but we do not support externalizing the broker.

--jmxuser admin

Specifies the username for shutting down the service.

Default: admin

--jmxpassword activemq

Specifies the password for shutting down the service.

Default: activemq

da_broker

Specifies the broker name that will be shut down.

Default: da_broker

3. Run the start command:

```
./activemq start
```

Chapter 3: Discovering Your Network

This section contains the following topics:

[Device Discovery](#) (see page 53)
[Discovery Workflow](#) (see page 54)
[SNMP Profiles](#) (see page 55)
[Discovery and Polling](#) (see page 57)
[Discovery and Polling in VMware Environments](#) (see page 59)
[Discovery Profiles](#) (see page 60)
[Run On Demand Discoveries](#) (see page 68)
[Schedule Discoveries](#) (see page 70)
[View Discovery Results](#) (see page 71)
[Discovery From Other Data Sources](#) (see page 73)
[Device Type Modifications](#) (see page 74)
[Rediscovery](#) (see page 76)

Device Discovery

Discovery is the process through which Data Aggregator discovers and models your IT infrastructure.

The discovery process does the following steps:

- Confirms what protocols devices respond to. Data Aggregator always determines if a device can respond to SNMP. If you select ICMP, Data Aggregator first determines if a device can respond to ICMP. If the device *does* respond to ICMP, Data Aggregator then determines if the device can respond to SNMP. If the device does *not* respond to ICMP, Data Aggregator will not confirm that the device responds to SNMP.
- Retrieves a minimum set of information for every discovered device that is sufficient to classify the device and add it to the appropriate device collection.

You can use two methods to discover devices in Data Aggregator:

- You can discover specific devices in your infrastructure environment using the discovery profiles that you create in Data Aggregator. [Follow the discovery workflow to manage devices with this method.](#) (see page 54)
- [You can discover devices that have been contributed from CA Performance Center](#) (see page 73).


Discovery Workflow

The following workflow offers a best practice to use as a quick reference when performing a discovery of your inventory.

Perform this process as either a user with the Administrator role or as the tenant administrator.

1. If you want Data Collector to perform queries of device MIB tables that use SNMP, create SNMP profiles in CA Performance Center before performing discovery.
Note: To apply an SNMP profile to a tenant, the SNMP profile must be *created* in the tenant space by a user with the tenant administrator role. For more information about creating SNMP profiles, see the *CA Performance Center Administrator Guide*.
2. [Create one or more discovery profiles from the Data Aggregator administration pages](#) (see page 61).
3. [Run one or more discovery profiles](#) (see page 68). Discovery can be scheduled or run manually.
4. [Review the discovery results](#) (see page 71).
5. [Review the component monitoring results](#) (see page 90). Use the results to determine how to manage the devices and components.

This diagram illustrates the discovery process:



More Information:

[Device Discovery](#) (see page 53)

SNMP Profiles

SNMP profiles are definitions that contain the information necessary to allow Data Collector to do queries of device MIB tables that use SNMP. Data Collector can communicate with devices that support SNMPv1, SNMPv2c, and SNMPv3. Community strings and credentials are encrypted when they are stored in CA Performance Center and when they are sent to Data Aggregator and Data Collector.

Important! When using SNMPv3 community names, CA Performance Management requires that any authentication passwords or privacy passwords are greater than eight characters in length. If they are configured with passwords that are shorter than eight characters in length, SNMPV3 profiles can be unsuccessful in communicating with devices.

Data Collector uses SNMPv1/SNMPv2c and SNMPv3 profiles during inventory discovery to determine what credentials to use when accessing a device. CA Performance Center maintains this list of profiles. Each profile is ranked for device access. During discovery, each profile is tried for device access. The profile with the highest rank that can access a device is used.

You can create SNMP profiles in CA Performance Center and you can change SNMP profiles rankings. A new ranked list of SNMP profiles takes effect in the following situations:

- A new device is discovered.
- An existing device becomes unreachable through SNMP for at least two poll cycles.
- The SNMP profile a device is using is deleted from CA Performance Center.

Otherwise, devices that are already being polled successfully continue to use the existing SNMP profile, regardless of any changes you made to the ranked list of SNMP profiles.

Note: If the SNMPv1/SNMPv2c profile is the highest ranked profile that can access a device and the device can be accessed with both SNMPv1 and SNMPv2c, Data Collector communicates with that device using SNMPv2c.

We tested Data Collector to determine the CPU load that is added when various SNMPv3 protocols are used. We found that SHA /AES had a moderate (< 30 percent) impact on CPU utilization as compared to SNMPv1. MD5/DES, SHA/DES, and SHA/3DES were found to have a major (>30 percent) impact on CPU utilization.

Note: The servers on which this testing was conducted have some AES capability that is built into the CPUs.

If you add additional CPU cores to your environment, Data Collector can balance the CPU load.

You create SNMP profiles in the CA Performance Center user interface or using the CA Performance Center REST web services. After you create the SNMP profiles, they synchronize immediately with Data Aggregator and are available for inventory discovery to use them.

Note: For more information about creating SNMP profiles, see the *CA Performance Center Administrator Guide* and the *CA Performance Center REST Web Services Guide*.

After you run a discovery, you can access the Discovery History view to see the list of SNMP profiles that are used and the highest ranked SNMP profile the device responded to.

More Information:

[View Discovery Results](#) (see page 71)

Discovery and Polling

Discovery is the process through which Data Aggregator discovers and models your IT infrastructure.

The discovery process does the following steps:

- Confirms what protocols devices respond to, depending upon what protocols you select when you create a discovery profile. For example, assuming that you select all protocols (SNMP and ICMP), the following steps occur. Data Aggregator determines if a device can respond to ICMP. Data Aggregator then determines if the device can respond to SNMP. If the device does *not* respond to ICMP, Data Aggregator will not confirm that the device responds to SNMP.
- Retrieves a minimum set of information for every discovered device that is sufficient to classify the device and add it to the appropriate device collection.

Inventory discovery is the process where Data Aggregator identifies the devices on your network. Devices are identified using the IP domain, IP addresses, IP ranges, and hostnames you specify in discovery profiles. Specifically, inventory discovery identifies if devices are manageable or not (pingable vs SNMP capable), and determines the classification (router, switch, and so on). Inventory discovery also determines the vendor (Cisco, Juniper, and so on) and determines the type (7700, 8200, and so on).

Devices that are discovered during this process are automatically added to out-of-the box device collections, depending on the rules that control each device collection membership. You can also create custom device collections in CA Performance Center that create corresponding custom device collections in Data Aggregator when synchronization occurs. During the first synchronization with CA Performance Center *after* devices are discovered, the devices are added to custom device collections per the rules that are defined on those device collections.

Note: For more information about creating custom device collections and synchronizing them with Data Aggregator, see the *CA Performance Center Administrator Guide*.

Component monitoring is a separate process. The monitoring process involves the collection and analysis of various operational data for specific device components, such as CPU, memory, and interfaces. All of the information that describes how the monitoring is done exists within monitoring profiles that you assign to device collections.

The relationship of monitoring profiles to device collections governs component monitoring. Component monitoring can be triggered in the following ways:

- A monitoring profile is assigned to a device collection which a given device is a member.
- A device is added to a device collection that already has a monitoring profile assigned to it.
- A monitoring profile, which is assigned to a device collection, is edited to include a new metric family to monitor. The components that are associated with the metric family are then monitored automatically for each device in the collection that the monitoring profile is attached to; if the component was not previously monitored for the device.
- A new vendor certification is added for an existing metric family that is polled in a monitoring profile.
- A monitoring profile specifies a change detection rate and has “Automatically Update Metric Families” checked.
- You click the “Update Metric Family” button on the Polled Metric Families tab of the Monitored Devices view.

A *metric family* defines the set of values to collect and report on for a given technology. These values are normalized so that reporting is uniform regardless of the data source. When included in a monitoring profile, metric families determine which values to collect for the devices that are associated with that monitoring profile.

Polling begins automatically after inventory discovery and component monitoring complete. Operational metrics and configuration data are polled on a discovered device and its monitored components. The operational metrics and configuration data that are polled depend on the metric families you specify in the monitoring profile. Operational metrics are collected and retained at regular intervals for reporting. Examples of operational metrics include error rate, daily baseline, hourly baseline, and port performance. Configuration data represents or identifies a component or the component configuration.

Examples of configuration data include:

ifNumber

A MIB variable that tells Data Aggregator how many ports a device has.

ifStackLastChange

A MIB variable that indicates whether a change occurs on an interface stack table.

Discovered devices and monitored components usually take up to 5 minutes to begin synchronizing with CA Performance Center. Devices and components that are discovered and monitored while a synchronization is in progress are synchronized after the current synchronization completes.

More information:

[View Monitored Devices](#) (see page 90)

[How to Manage Change Detection](#) (see page 100)

Discovery and Polling in VMware Environments

You can discover and monitor your VMware virtual machines and ESX hosts alongside your network devices. Although your VMware devices and components behave like physical components, the process for discovering and monitoring those devices and components differs to accommodate the collection of data from vCenter. Although you can discover the VMs and ESX Hosts directly with SNMP, you may also want to collect vCenter data through the vCenter Server Application Insight Module (VCAIM).

You can discover ESX hosts and virtual machines in your VMware environment.

During inventory discovery, Data Aggregator identifies ESX hosts and virtual machines in the following ways:

- Through ICMP
- Through SNMP if the servers have an SNMP agent deployed
- Through discovery of a server running systemEdge with the VCAIM

Although each ESX host and virtual machine can be identified multiple times through ICMP, SNMP, and vCenter, only one device is created. This device represents the ESX host or virtual machine.

Once the ESX and VM devices are created, Data Aggregator begins polling any vCenter specific metrics, and can discover and begin polling more components that are identified through the SNMP agent.

Depending on the source of the metric data, some polling for the VMs and ESX is done by direct polling on the device, while the VCAIM is polled to collect other data.

By default, every 15 minutes, after Data Aggregator discovers your VMware environment, it monitors for virtual machines that have been added or removed, or VMotioned between ESX hosts. By default, every 24 hours, Data Aggregator also monitors for ESX hosts that have been added or removed.

Discovery Profiles

Discovery profiles specify how inventory discovery operates. As an administrator, you can use the CA Performance Center user interface or Data Aggregator REST web services to manage discovery profiles.

Within a discovery profile, you specify the IP addresses, IP address ranges, and host names for which you want to discover devices. You also specify an IP domain. You can only specify one IP domain for each discovery profile you create. Newly discovered devices will be created within that IP domain.

When multiple Data Collector hosts are deployed in one IP domain, each Data Collector issues a discovery request to that device.

When more than one Data Collector can contact the same device, a specific Data Collector is selected to monitor the device. An algorithm, which is based on load balancing, determines this selection.

IP domains are also necessary for monitoring tenant environments with overlapping IP addresses. One tenant can have one or multiple IP domains. If a tenant has overlapping IP addresses, there must be multiple IP domains in the network. Overlapping IP addresses are handled through IP domains.

IP domains are created in CA Performance Center. Data Aggregator becomes aware of new IP domains when manual or automatic synchronization occurs.

The discovery process attempts to distribute the devices across the available Data Collector instances, but this process does not take into account which devices a Data Collector instance is currently monitoring.

Note: For more information about creating and synchronizing IP domains, see the *CA Performance Center Administrator Guide*.

Discovery profiles are only accessible by users within the tenant space that the discovery profile was created in. A user that is assigned to the 'Default Tenant' space can run a discovery using a discovery profile that exists in the 'Default Tenant' space and can see the results of that discovery.

Therefore, it is important to be logged in as, or be administering, the correct tenant *before* you create a discovery profile.

Note: For more information about creating and administering tenants, see the *CA Performance Center Administrator Guide*.

More information:

[Create Discovery Profiles](#) (see page 61)

[Delete Discovery Profiles](#) (see page 68)

[Edit Discovery Profiles](#) (see page 65)

[View a List of Discovery Profiles](#) (see page 61)

[IP Domain Deletions](#) (see page 96)

View a List of Discovery Profiles

Discovery profiles for SNMP and ICMP let you configure how discovery operates in your environment.

You can view a list of discovery profiles and the details of each. You can view the status of discovery and the time when discovery was last run. These details help you understand how your network is being discovered.

Note: Log in as a tenant administrator to perform this task.

Follow these steps:

1. Select Admin, Data Source Settings, and click a Data Aggregator data source.
2. Click Discovery Profiles from the Monitored Inventory menu.

The Discovery Profiles List page opens, displaying a list of available discovery profiles.

More information:

[Discovery Profile IP Ranges](#) (see page 64)

[Create Discovery Profiles](#) (see page 61)

[Delete Discovery Profiles](#) (see page 68)

[Edit Discovery Profiles](#) (see page 65)

[Run On Demand Discoveries](#) (see page 68)

[Schedule Discoveries](#) (see page 70)

[Discovery Profiles](#) (see page 60)

[View Discovery Results](#) (see page 71)

Create Discovery Profiles

You can create discovery profiles to specify how inventory discovery operates in your environment.

Log in as a tenant administrator to perform this task. Discovery profiles are only accessible by users within the tenant space that the discovery profile was created in.

Note: For more information about tenants, see the *CA Performance Center Administrator Guide*.

Follow these steps:

1. [Navigate to the list of available discovery profiles](#) (see page 61).
2. Click New.
3. Do the following steps:
 - a. Type a descriptive name for the discovery profile in the Name field.

Note: Single quotes, double quotes, backward slashes, forward slashes, and ampersands are not permitted.
 - b. Select an IP domain from the list of pre-configured domains.
4. Select the IPs/Hosts tab and do one or more of the following actions:
 - (Optional) Navigate to and import a CSV file of IP addresses. The CSV file can contain a comma-separated list of IPv4 addresses, IPv6 addresses, IPv4 address ranges, and hostnames. Browse to select the file and click Open.

Note: For Chinese characters to be applied to the alias name, save the CSV file in UTF-8 format.
 - Type IP address ranges for which you want to discover devices in the IP Address Range field. Comma-delimited values are accepted.

Note: If an IP range includes multiple IP addresses from a device which has a hostname and the IP that maps to the hostname is also included in the IP range, inventory discovery will always use the hostname IP for the device's primary IP address.
 - Type individual IP addresses for which you want to discover devices in the IP Address List field. Comma-delimited values are accepted.
 - Type the host names for which you want to discover devices in the Host List field. Comma-delimited values are accepted.
 - Copy a list of individual IP addresses, IP address ranges, and hostnames to the clipboard, and then paste the list into the list view by pressing ctrl+v.
 - Remove an item from the IP list by selecting the IP address, range, or hostname, and clicking Delete.
 - Search for an item in the IP list by entering the IP address, range, or hostname in the Search field. To return to the full list of items in the IP list, click the X button. Alternatively, you can press the Esc button on your keyboard.

Note: Edit an IP address, range, or hostname in the IP list by double-clicking it. Press Enter to save the changes. Press Esc to exit edit mode without saving the changes.

Do not include duplicate IP addresses or hostnames. If duplicates are detected, a message displays, indicating that duplicates were found and ignored.

Note: Single quotes, double quotes, backward slashes, forward slashes, and ampersands are not permitted.

5. (Optional) Select the Schedule tab. To create a schedule for when you want this discovery profile to run, do the following steps:
 - To create a daily schedule, select Daily from the Scheduling Interval drop-down box. Select the time when you want the discovery to begin each day.
 - To create a weekly schedule, select Weekly from the Scheduling interval drop-down box. Select each day that you want the discovery to run. Select a time when you want the discovery to begin.
6. Select the SNMP tab. If you want to use all SNMP profiles, you do not need to do anything. All SNMP profiles are selected by default. To use specific SNMP profiles, select the 'Use specific list of assigned SNMP profiles checkbox'. Select one or more SNMP profiles from list of available profiles and move them to the assigned list. Using a subset of SNMP profiles can help reduce network traffic.
7. Select the Advanced tab and do the following steps:
 - a. (Optional) Change the priority in which you want the discovered device to be named. During discovery, any device item that the discovery profile creates is named with the highest available naming convention. If a naming convention is not set in the MIB for the device, it is not available and the next highest priority naming convention is tried.
 - b. (Optional) Select the 'Save as Default' option if you want to save the naming order for *new discovery profiles*. The next time that you create a discovery profile, the naming order automatically appears in the order you saved.

The out-of-the-box default naming order is System Name, Host Name, IP Address.
 - c. Select 'Use ICMP' if you want Data Aggregator to determine if a device can respond to ICMP during the discovery process. Select 'Create Pingables' to create pingable devices during discovery. Deselect 'Use ICMP' to prevent the creation of pingable devices. Data Aggregator does not determine if a device can respond to ICMP when these options are deselected.

Select the 'Save as Default' option if you want to save the ICMP discovery options you chose. The next time that you create a discovery profile, the ICMP discovery options are automatically selected.
8. Click Save.

The discovery profile is created and is displayed in the Discovery Profiles list.

More information:

[Discovery Profile IP Ranges](#) (see page 64)

[Discovery Workflow](#) (see page 54)

[Discovery Profiles](#) (see page 60)

Discovery Profile IP Ranges

When you create or edit a discovery profile, you can enter the IP address ranges you want to discover for IPv4. Range discovery is not supported for IPv6 addresses.

When specifying IP ranges in the discovery profile, the following rules apply:

- An IPv4 range can contain wildcards (*). A wildcard represents a full range for an IP octet: 0-255.
- An IPv4 range can contain hyphens (-). A hyphen can exist between the lower IP address and upper IP address. A hyphen can also be in the IP octets in the lower IP address.
- If a wildcard or a hyphen is used in a term in the lower IP address, the upper IP address cannot be represented.

Examples: Valid IP Ranges

- Both of the following examples attempt to discover devices at every IP address from 10.25.1.0 to 10.25.1.190:

10.25.1.0-10.25.1.190

OR

10.25.1.0-190

- Both of the following examples attempt to discover devices at every IP address from 10.25.0.0 to 10.25.255.255:

10.25.*.*

OR

10.25.0.0 - 10.25.255.255

- Both of the following examples attempt to discover devices at every IP address from 10.25.0.3 to 10.25.0.40 and from 10.25.1.3 to 10.25.1.40:

10.25.0-1.3-40

OR

10.25.0.3 - 10.25.0.40, 10.25.1.3 - 10.25.1.40

- Both of the following examples attempt to discover devices at every IP address from 10.25.0.0 to 10.25.0.5, from 10.25.1.0 to 10.25.1.5, and so on, up to 10.25.255.0 to 10.25.255.5:

10.25.*.0-5

OR

10.25.0.0 - 10.25.0.5, 10.25.1.0 - 10.25.1.5 ... 10.25.255.0 - 10.25.255.5

Examples: Invalid IP Ranges

- The following example is invalid because the upper IP address is incomplete:

10.25.1.0 - 10.23

- The following example is invalid because when a hyphen (-) is used in an octet in the lower IP address, the upper IP address cannot be present:

10.25.1.0-190 - 10.25.1.255

- The following example is invalid because when a wildcard (*) is used in an octet in the lower IP address, the upper IP address cannot be present:

10.25.*.0 - 10.25.255.255

- The following example is invalid because it is unclear whether the wildcard octet (1*) implies 10.25.10-19.0 or 10.25.10-199.0:

10.25.1*.0

More information:

[Create Discovery Profiles](#) (see page 61)

[View a List of Discovery Profiles](#) (see page 61)

Edit Discovery Profiles

You can edit an existing discovery profile.

Note: Log in as a tenant administrator to perform this task.

Follow these steps:

1. [Navigate to the list of available discovery profiles](#) (see page 61).

2. Select the discovery profile that you want to edit and click Edit. Modify the different fields on each tab as needed.
3. So the following steps:
 - a. Type a descriptive name for the discovery profile in the Name field.
Note: Single quotes, double quotes, backward slashes, forward slashes, and ampersands are not permitted.
 - b. Select an IP domain from the list of pre-configured domains.
Note: If you already ran a discovery on this discovery profile, you cannot change the IP domain.
4. Select the IPs/Hosts tab and do one or more of the following actions:
 - (Optional) Navigate to and import a CSV file of IP addresses. The CSV file can contain a comma-separated list of IPv4 addresses, IPv6 addresses, IPv4 address ranges, and hostnames. Browse to select the file and click Open.
Note: For Chinese characters to be applied to the alias name, save the CSV file in UTF-8 format.
 - Type IP address ranges for which you want to discover devices in the IP Address Range field. Comma-delimited values are accepted.
Note: If an IP range includes multiple IP addresses from a device which has a hostname and the IP that maps to the hostname is also included in the IP range, inventory discovery will always use the hostname IP for the device's primary IP address.
 - Type individual IP addresses for which you want to discover devices in the IP Address List field. Comma-delimited values are accepted.
 - Type the host names for which you want to discover devices in the Host List field. Comma-delimited values are accepted.
 - Copy a list of individual IP addresses, IP address ranges, and hostnames to the clipboard, and then paste the list into the list view by pressing ctrl+v.
 - Remove an item from the IP list by selecting the IP address, range, or hostname, and clicking Delete.
 - Search for an item in the IP list by entering the IP address, range, or hostname in the Search field. To return to the full list of items in the IP list, click the X button. Alternatively, you can press the Esc button on your keyboard.

Note: Edit an IP address, range, or hostname in the IP list by double-clicking it. Press Enter to save the changes. Press Esc to exit edit mode without saving the changes.

Do not include duplicate IP addresses or hostnames. If duplicates are detected, a message displays, indicating that duplicates were found and ignored.

Note: Single quotes, double quotes, backward slashes, forward slashes, and ampersands are not permitted.

5. (Optional) Select the Schedule tab. To create a schedule for when you want this discovery profile to run, do the following steps:
 - To create a daily schedule, select Daily from the Scheduling Interval drop-down box. Select the time when you want the discovery to begin each day.
 - To create a weekly schedule, select Weekly from the Scheduling interval drop-down box. Select each day that you want the discovery to run. Select a time when you want the discovery to begin.
6. Select the SNMP tab. If you want to use all SNMP profiles, you do not need to do anything. All SNMP profiles are selected by default. To use specific SNMP profiles, select the 'Use specific list of assigned SNMP profiles checkbox'. Select one or more SNMP profiles from list of available profiles and move them to the assigned list. Using a subset of SNMP profiles can help reduce network traffic.
7. Select the Advanced tab and do the following steps:
 - a. (Optional) Change the priority in which you want the discovered device to be named. During discovery, any device item that the discovery profile creates is named with the highest available naming convention. If a naming convention is not set in the MIB for the device, it is not available and the next highest priority naming convention is tried.
 - b. (Optional) Select the 'Save as Default' option if you want to save the naming order for *new discovery profiles*. The next time that you create a discovery profile, the naming order automatically appears in the order you saved.

The out-of-the-box default naming order is System Name, Host Name, IP Address.
 - c. Select 'Use ICMP' if you want Data Aggregator to determine if a device can respond to ICMP during the discovery process. Select 'Create Pingables' to create pingable devices during discovery. Deselect 'Use ICMP' to prevent the creation of pingable devices. Data Aggregator does not determine if a device can respond to ICMP when these options are deselected.

Select the 'Save as Default' option if you want to save the ICMP discovery options you chose. The next time that you create a discovery profile, the ICMP discovery options are automatically selected.
8. Click Save.

The discovery profile is updated with your changes. The next time this discovery profile runs, your changes are applied.

More information:

[Troubleshooting: Discovery Does Not Start](#) (see page 153)
[Discovery Profiles](#) (see page 60)

Delete Discovery Profiles

If you no longer want a discovery profile, you can delete it. For example, you can delete discovery profiles that are no longer used, are duplicates of other discovery profiles, and so on. You cannot rediscover the devices that are specified in a deleted discovery profile.

Note: Log in as a tenant administrator to perform this task.

Follow these steps:

1. [Navigate to the list of discovery profiles](#) (see page 61).
2. Select the discovery profile that you want to delete and click Delete.
A confirmation dialog opens.
3. Click Yes.
The discovery profile is deleted.

More Information:

[Discovery Profiles](#) (see page 60)

Run On Demand Discoveries

Inventory discovery is the process through which devices are discovered in your network, which is based on information you add to discovery profiles. You can run an on-demand discovery.

Attempts to discover devices are made through SNMP and ICMP protocols. If a device does not respond to SNMP with the SNMPv1/SNMPv2c or SNMPv3 profiles you created but *does* respond to ICMP, a pingable device is created. (SNMP profiles are created using either the CA Performance Center user interface or the CA Performance Center REST web services.)

This task can be done as either the tenant administrator or as the administrator. To run a discovery as the administrator, configure Data Collector for the Default Tenant domain before running the discovery.

Note: For more information about creating SNMP profiles and synchronizing them with Data Aggregator, see the *CA Performance Center Administrator Guide* and the *CA Performance Center REST Web Services Guide*.

Follow these steps:

1. [Navigate to the list of discovery profiles in CA Performance Center](#) (see page 61).

2. Select one or more discovery profiles that you want to run a discovery on and click Run.

Note: You can only run a discovery on a discovery profile that has a state of 'READY'.

A confirmation dialog opens.

3. Click Yes.

The discovery starts, the State column for the selected discovery profiles indicates 'RUNNING', and the Last Run Time column is updated with the time when the discovery starts.

Note: To trigger the Percentage Complete column to update while discovery is running, click Refresh.

A confirmation dialog opens.

4. Click OK.

Discovered devices are added to device collections, which initiates component monitoring and polling. You are returned to the Discovery Profile page.

If discovery hangs for more than 10 minutes, it is aborted. A discovery is considered to be hanging when no new devices are discovered within 10 minutes *and* the state for the selected discovery profiles have not changed within 10 minutes. An audit event is generated on the Data Aggregator device.

The State column for the selected discovery profiles indicates 'FAILURE' if no devices were discovered successfully, or 'PARTIAL_FAILURE' if at least one device was discovered successfully.

The discovered devices and monitored components can take up to 5 minutes to synchronize with CA Performance Center. When the synchronization is complete, the discovered devices and monitored components appear in the Inventory tab in CA Performance Center.

5. (Optional) To synchronize discovered devices and monitored components immediately with CA Performance Center, perform the following steps:
 - a. Select Admin, Data Source Settings, and click a Data Aggregator data source.
 - b. Click Data Aggregator from the System Status menu.
 - c. Select Data Aggregator and click the Resynch button.

More information:

[Discovery Workflow](#) (see page 54)

[How to Set and Activate an Interface Filter](#) (see page 115)

Schedule Discoveries

Inventory discovery is the process through which devices are discovered in your network, which is based on information you add to discovery profiles. You can schedule a discovery to run daily or weekly.

Note: You can run a scheduled discovery at any time by selecting a discovery profile and clicking Run. However, you cannot initiate an on-demand discovery for a discovery profile while a scheduled discovery for the discovery profile is in progress.

Attempts to discover devices are made through SNMP and ICMP protocols. If a device does not respond to SNMP with the SNMPv1/SNMPv2c or SNMPv3 profiles you created but *does* respond to ICMP, a pingable device is created. (SNMP profiles are created using either the CA Performance Center user interface or the CA Performance Center REST web services.)

To run a discovery as the administrator, configure Data Collector for the Default Tenant domain before scheduling the discovery.

Note: For more information about creating SNMP profiles and synchronizing them with Data Aggregator, see the *CA Performance Center Administrator Guide* and the *CA Performance Center REST Web Services Guide*.

Follow these steps:

1. [Navigate to the list of discovery profiles in CA Performance Center](#) (see page 61).
2. Do one of the following steps:
 - Select an existing discovery profile that you want to schedule a discovery for and click Edit.
The Edit Discovery Profile page opens.
 - Click New to create a discovery profile that you want to schedule a discovery for.
The Discovery Profiles dialog opens.
3. To create a schedule for when you want this discovery profile to run, do one of the following steps:
 - To create a daily schedule, select Run Daily from the Scheduling Interval drop-down box and then select the time when you want the Discovery to begin each day.
 - To create a weekly schedule, select Run Weekly from the Scheduling interval drop-down box, select each day when you want the discovery to run, and select a time when you want the discovery to begin.

Note: Select None from the Scheduling drop-down box to remove a schedule.
4. Click Save.

When the discovery is scheduled, the State column for the discovery profile indicates 'SCHEDULED' and the next scheduled run time is displayed.

When the scheduled discovery starts, the State column for the selected discovery profiles indicates 'RUNNING' and the Last Run Time column is updated with the time when the discovery starts.

Note: To trigger the Percentage Complete column to update while discovery is running, click Refresh.

Discovered devices are added to device collections, which initiates component monitoring and polling. You are returned to the Discovery Profile page.

If discovery hangs for more than 10 minutes, it is aborted. A discovery is considered to be hanging when no new devices are discovered within 10 minutes *and* the state for the selected discovery profiles have not changed within 10 minutes. An audit event is generated on the Data Aggregator device.

The State column for the selected discovery profiles indicates 'FAILURE' if no devices were discovered successfully, or 'PARTIAL_FAILURE' if at least one device was discovered successfully.

The discovered devices and monitored components can take up to 5 minutes to begin synchronizing with CA Performance Center. When the synchronization is complete, the discovered devices and components appear in the Inventory tab in CA Performance Center.

5. (Optional) To synchronize discovered devices and components immediately with CA Performance Center, perform the following steps:
 - a. Select Admin, Data Source Settings, and click a Data Aggregator data source.
 - b. Click Data Aggregator from the System Status menu.
 - c. Select Data Aggregator and click the Resynch button.

View Discovery Results

You can view a summary of the number of new pingable (ICMP) and manageable (SNMP) devices that were discovered during a specific discovery instance. You can also view specific details about these discovered devices, including the IP address, model, type, vendor name, location, and the protocols that were used.

A discovery result can also indicate that existing devices were found. The same or a different discovery profile discovered these existing devices previously. To view existing devices, use the Unchanged filter. Existing devices that show different IP addresses indicate that the devices were previously discovered and are being monitored with a different IP address.

This behavior is common and expected, as many devices can respond to multiple IP addresses. Data Aggregator maintains the full set of IP addresses for each device and contributes them all to CA Performance Center.

Follow these steps:

1. [View a list of discovery profiles](#) (see page 61).
2. Select a discovery profile for which you want to view discovery results and click the History button.
Note: The History button is disabled when you select a discovery profile for which a discovery has not been run.
3. Select a Discovery instance, if applicable.
4. (Optional) Filter the Discovered Devices table by doing one of the following options:
 - Filter by device type by selecting which device types you want displayed from the Device Type Filter list and clicking Apply.
 - Filter by the state of the discovered devices by selecting which states you want displayed from the State list and clicking Apply.
 - Filter by device type and state by selecting from the Device Type Filter list and the State list and clicking Apply.

The discovery results appear in the Discovered Devices table. The SNMP Profile column shows the highest ranked SNMP profile the device responded to.

Specifically, the State column indicates one of the following states:

New

Indicates a device that was discovered for the first time when this discovery profile was run.

Changed

Indicates that a device type has changed from a previous discovery. For example, a previously discovered pingable device is now discovered as a manageable device. Or a previously manageable device with the device type of “Switch” has now changed to the device type of “Router”, and so on. Devices with only attribute changes, such as hostname, system description, and so on, are not classified as Changed.

Unchanged

Indicates that existing devices have not changed. Existing devices with only attribute changes are classified as Unchanged also.

Deleted

Indicates that the device was deleted from Data Aggregator since a discovery was run.

Note: If an individual discovered device fails to be recognized as pingable or manageable, its state indicates Unreachable. However, Data Aggregator does *not* report unreachable devices that are found in an IP range.

More information:

[SNMP Profiles](#) (see page 55)

[Discovery Workflow](#) (see page 54)

Discovery From Other Data Sources

You can choose whether Data Aggregator automatically discovers devices that are synchronized with CA Performance Center by other data sources. This option is available when you register Data Aggregator or when you edit data source options. By default, this option is disabled.

Important! When enabled, Data Aggregator attempts to discover devices that are contributed by *all* other data sources. You cannot refine this capability to a particular set of data sources.

When enabled, Data Aggregator attempts to discover any new devices it learns about from that point forward. If you want Data Aggregator to attempt to discover devices that were synchronized with CA Performance Center in the past, select the Data Aggregator data source, click Resync, and select the Perform a Full Resynchronization check box.

The discovery attempt produces a pingable or other type of device in Data Aggregator if the device can be reached through ICMP or some other supported protocol.

Note: If you disable this option anytime after it is enabled, Data Aggregator continues to monitor any devices that were discovered already.

To enable this option, select the Discover devices from other data sources checkbox on the Edit Data Source dialog on the Manage Data Sources page in CA Performance Center.

More information:

[Device Discovery](#) (see page 53)

Device Type Modifications

Based on the device service information, Data Aggregator can automatically classify manageable devices as Router, Switch, and Server types. If a manageable device cannot be identified as a Router, a Switch or a Server, it is classified as the 'Device' device type.

If the types of some SNMP manageable devices were not identified as you expected, you can override the device types. Map the device sysObjectID MIB value explicitly to the correct device type in the \$KARAF_HOME/custom/devicetypes/DeviceTypes.xml file that is shipped with Data Aggregator.

Note: You cannot add new device types to the DeviceTypes.xml file.

The DeviceTypes.xml file contains a template to map the sysObjectID to appropriate device types. By default, the file does not contain any sysObjectID-to-type mapping entry. If you want to classify a device type with a particular sysObjectID, you can modify the template to add the sysObjectID-to-type entries into the file. Before you add a sysObjectID, uncomment the section where you are adding the sysObjectID.

Note: Updates to the DeviceTypes.xml file can take up to one minute to apply.

A device can be classified into multiple device types. However, the type, Device, is mutually exclusive to other device types. For example, if you add a sysObjectID to one or more of the Router, Switch, or Server device types and you also add that sysObjectID to the 'Device' device type, the 'Device' device type is dropped and is not recognized.

Note: If you upgrade Data Aggregator, the DeviceTypes.xml file is not preserved. However, the configurations that were added prior to the upgrade *are* preserved.

Example: Map a Device sysObjectID to Another Device Type

Follow these steps:

1. Open the \$KARAF_HOME/custom/devicetypes/DeviceTypes.xml file.

2. Enter the following information:

```
<DeviceType>
  <Routers>
    <sysObjectID>1.3.6.5.1.34</sysObjectID>
  </Routers>

  <Switches>
    <sysObjectID>1.3.6.5.5.3</sysObjectID>
    <sysObjectID>1.3.6.5.1.34</sysObjectID>
  </Switches>

  <Servers>
    <sysObjectID>1.3.6.5.567.1</sysObjectID>
  </Servers>

  <Device>
    <sysObjectID>1.3.6.5.49.1</sysObjectID>
  </Device>
</DeviceType>
```

3. Run discovery on the discovery profile that contains the devices.

Note: The changes that you make to the DeviceTypes.xml file do not take effect on existing devices until you rerun discovery.

When discovery is run, the following results occur:

- All devices that have a sysObjectID of 1.3.6.5.1.34 are classified as a device type of Router and Switch.
- All devices that have a sysObjectID of 1.3.6.5.5.3 are classified as a device type of Switch.
- All devices that have a sysObjectID of 1.3.6.5.567.1 are classified as a device type of Server.
- All devices that have a sysObjectID of 1.3.6.5.49.1 are classified as a device type of Device.

Rediscovery

Existing monitored devices can be rediscovered when a discovery profile containing one of its IP addresses or its host name is run. A single monitored device can be rediscovered when you click the Rediscover button on the Details tab for the specific device.

The following set of attributes can be updated as a result of this discovery:

- System name
- Hostname
- Device type (as it appears in CA Performance Center)
- Location
- Vendor
- Device description
- Device model

Note: Changes to device attributes can result in changes to the groups and device collections a device is a part of. Changes to groups and device collections can potentially add or remove monitoring profiles.

Changes in the device attributes can take up to 5 minutes to be seen in CA Performance Center inventory or dashboards views.

Chapter 4: Managing the Infrastructure

This section contains the following topics:

[Customizing Device and Component Management Workflow](#) (see page 77)
[Monitoring Profiles](#) (see page 80)
[Factory Device Collections](#) (see page 85)
[Custom Device Collections](#) (see page 89)
[View Monitored Devices](#) (see page 90)
[Delete Devices](#) (see page 92)
[Change the Primary IP Address for a Monitored Device](#) (see page 93)
[Delete Retired Components](#) (see page 94)
[IP Domain Deletions](#) (see page 96)
[Tenant Deletions](#) (see page 97)
[Disable Tenants](#) (see page 98)
[Enable Tenants](#) (see page 99)
[Device Reconfiguration](#) (see page 99)

Customizing Device and Component Management Workflow

You can customize the management of your discovered devices and monitored components. Options include modifying profiles, modifying associations, creating new vendor certifications, and importing metric families. For example, you can poll critical interfaces more frequently or can apply custom monitoring profiles with event rules to custom device collections.

The following workflow offers a best practice to use as a quick reference for customization.

Log in as a user with the Administrator role and perform the following steps:

1. Create new monitoring profiles (or make copies of factory monitoring profiles) to customize the poll rates and metrics for monitoring your devices.
2. (Optional) [Add event rules to custom monitoring profiles](#) (see page 128).
3. (Optional) If the factory vendor certifications and their associated metric families do not meet your needs, create custom vendor certifications and import new metric families. This step can be done at any time.

Note: For more information about custom metric families and custom vendor certifications, see the *Data Aggregator Self-Certification Guide*.

4. Create custom device collections and associated rules in CA Performance Center, which are then used as Data Aggregator device collections. You can either synchronize these device collections with Data Aggregator immediately or wait for the automatic synchronization to occur. The population of devices into these device collections can be done manually after discovery.

Note: If you are an MSP or a tenant, perform this step as the tenant administrator. For more information about creating monitored groups and synchronizing data sources, see the *CA Performance Center Administrator Guide*.

5. [Customize your monitoring profile and device collection associations to make sure that the poll rate you want is used](#) (see page 83). When you create a custom monitoring profile, associate the custom monitoring profile with a custom device collection to activate the monitoring profile and any related event rules.

Note: If you are an MSP or a tenant, perform this step as the tenant administrator.

Customization can also include removing associations between factory monitoring profiles and device collections, and associating custom monitoring profiles to either factory or custom device collections.


6. [Review the component monitoring results after polling with the new configuration has begun to verify that you are collecting the information that you want](#) (see page 90).

Note: If you are an MSP or a tenant, perform this step as the tenant administrator.


This diagram illustrates the workflow for an enterprise:

Manage Devices and Components in an Enterprise Environment

Administrator



This diagram illustrates the workflow for a tenant environment:



Monitoring Profiles

Monitoring profiles determine the polling speed and what statistics are discovered and polled for the devices in a device collection. A selection of factory (out-of-the-box) monitoring profiles is provided. Factory monitoring profiles are automatically applied to factory device collections, such as the All Routers device collection. Factory monitoring profiles cannot be edited or deleted from the system, but they can be removed from a device collection or copied to create a custom profile.

You can create, edit, copy, or delete customized monitoring profiles as the administrator. Custom monitoring profiles are made available globally in the user interface and are not tenant-scoped even if you are an MSP administrator working in a tenant workspace. (However, the device collections that monitoring profiles are associated with are tenant-scoped). For example, you can create a "Gold Service Router Monitoring" monitoring profile and you can use it for all of your gold level tenants. You do not have to create a *separate* "Gold Service Router Monitoring" monitoring profile for *each* gold level tenant.

Name uniqueness of the monitoring profiles is enforced across all tenants.

Use the CA Performance Center user interface or Data Aggregator REST web services to manage monitoring profiles and view their associations with device collections.

Factory Monitoring Profile Associations

Monitoring profiles specify the statistics to poll. Factory (out-of-the-box) monitoring profiles are automatically associated with device collections as follows:

- The Accessibility monitoring profile is associated with the All Devices device collection.
- The Reachability monitoring profile is associated with the All Devices device collection.
- The Router monitoring profile is associated with the All Routers device collection.
- The Physical Server monitoring profile is associated with the All Servers device collection.
- The Virtual Server monitoring profile is associated with the All Servers device collection.
- The Switch monitoring profile is associated with the All Switches device collection.
- The Microsoft Cluster Services monitoring profile is associated with the All Servers device collection.
- The VMWare monitoring profile is associated with the All VMare vCenters device collection.
- The VMware ESX Host monitoring profile with the All VMware vCenters device collection.
- The VMware Virtual Machine monitoring profile with the All VMware vCenters device collection.

The following monitoring profiles do not have factory associations with device collections. This design prevents large discoveries that can impact performance. Manually assign these monitoring profiles to device collections to collect data for them:

- Network Interface
- Response Path
- MPLS
- CBQoS

More information:

[Factory Device Collections](#) (see page 85)

[All Devices Device Collection](#) (see page 87)

[All Routers Device Collection](#) (see page 87)

[Assign or Remove Monitoring Profiles from Device Collections](#) (see page 83)

View Monitoring Profiles

Administrators can view a list of monitoring profiles and their associations with device collections:

- Administrators can view the device collections for the tenant they are administering.
- A tenant administrator can view its own list of device collections.

This information helps you determine how to manage your monitoring profiles and poll rates and gives an idea of what report types can be produced for a device collection.

Follow these steps:

1. Click Monitoring Profiles from the Monitoring Configuration menu for a Data Aggregator data source.
A list of monitoring profiles is populated.
2. You can add a monitoring profile to the system, or you can select a monitoring profile to edit, copy, or delete if you are an administrator. All monitoring profiles, including custom, are global.

Note: Factory monitoring profiles cannot be edited or deleted; only custom monitoring profiles can be modified.

3. Select a monitoring profile.

4. Details for the selected monitoring profile populate the tabs, as follows:
 - The Metric Families tab is populated with a list of metric families that are associated with that specific monitoring profile. Metric families contain the metrics that are used for polling devices and interfaces.
 - The Event Rules tab is populated with a list of event rules that are associated with that specific monitoring profile. As the administrator, you can manage the relationships between event rules and the selected monitoring profile by assigning or removing rules.
 - The Collections tab is populated with a list of device collections that are associated with that specific monitoring profile. As a tenant administrator, you can manage the relationships between a device collection and the selected monitoring profile by assigning or removing profiles.

More information:

[Troubleshooting: Polling Has Stopped on Discovered Metric Family](#) (see page 154)
[Update Device Reconfiguration Automatically](#) (see page 102)

Assign or Remove Monitoring Profiles from Device Collections

Administrators or tenant administrators can add or remove a relationship between a specific device collection and the monitoring profiles in the system. This ability lets you start or stop polling the statistics that are associated with a monitoring profile in relation to the devices and components in a device collection.

Important! When you assign monitoring profiles to device collections, significant SNMP requests can occur. These requests can impact the device performance. Also, do not associate monitoring profiles with the All Devices device collection. Doing so can cause extra SNMP requests being made to pingable-only devices, and can result in sporadic metric family support.

For example, you have a discovered router with 1000 physical and logical interfaces. You also created an interface monitoring profile and you set the poll rate on that monitoring profile to one minute. If the interface monitoring profile is assigned to the device collection that your router is in, ten MIB objects are polled for each interface. This setting induces a rate of 166 MIB objects per second for the SNMP agent to reply to. This significant SNMP load can impact the performance of the router.

Metric families such as QoS, MPLS, and IPSLA can also contribute to significant SNMP requests. For more information on the implications and restrictions of SNMP requests to your network devices, refer to the vendor manuals or contact the vendor.

Note: Data Aggregator uses the fastest poll rate that is configured when multiple device collections are assigned to a monitoring profile. Remove the factory monitoring profile associations with device collections when you want your custom poll rate used.

Follow these steps:

1. Click Collections from the Monitoring Configuration menu for a Data Aggregator data source.

A list of collections appears. Administrators can view the device collections for the tenant they are administering. A tenant administrator can view its own (tenant) list of device collections.

2. Select a collection and click the Monitoring Profiles tab.

A list appears that shows monitoring profiles that are assigned to the selected device collection.

3. Click Manage.

The Assign Collection Monitoring Profiles dialog opens.

4. Do any of the following actions:

- Select one or more monitoring profiles from the Available Monitoring Profiles list, and click Add.

The selected monitoring profiles move to the Assigned Monitoring Profiles list.

Associating a monitoring profile to a device collection activates the event rules that are included in the monitoring profile. Events are raised and cleared when the devices in the device collection satisfy event rule conditions.

- Select one or more monitoring profiles from the Assigned Monitoring Profiles list, and click Remove.

The selected monitoring profiles move to the Available Monitoring Profiles list.

Note: Removing the relationship does not remove the monitoring profiles from the system.

5. Click Save.

Your changes are saved, and can be verified by repeating step 2.

More information:

[Factory Monitoring Profile Associations](#) (see page 81)

[Troubleshooting: Polling Has Stopped on Discovered Metric Family](#) (see page 154)

[Customizing Device and Component Management Workflow](#) (see page 77)

[How to Set and Activate an Interface Filter](#) (see page 115)

caim--Configure a Monitoring Profile Poll Filter

Filtering specifies the component items to be polled, and the time interval at which they are polled. Specifying which component items to poll lets you monitor only those items that interest you. For your *custom* monitoring profile, you can specify additional filters.

You can add or edit a filter before or after you run a discovery. Data Aggregator applies filtering after discovery. Only the component items that match the filter criteria are polled. If you add or edit a filter *after* you run a discovery, polling on these component items stops.

Note: Log in as the administrator to perform this task.

Follow these steps:

1. Select a Monitoring Profile that you have created in the list.
The details for the selected monitoring profile appear in the pane on the right. The metric family tab is selected by default.
2. Click the name of a metric family in the list.
The Edit Filter and Clear Filter buttons at the bottom of the pane become available.
3. Click the Edit Filter button.
The Filter Expression dialog appears.
4. Click the existing AND condition, and then click a logic button on the right-hand side of the dialog.
5. Select an attribute and operation, and enter a value for your condition.
6. Click the Add Condition button.
The condition that you have created is added to the filter expression.
7. Create any additional conditions.
Add each condition by clicking the Add Condition button.
8. Click the Save button. The filter expression is saved and assigned to the selected Metric Family.

Note: When viewing component items, and the filters assigned to them, an asterisk (*) appears next to each component item that does not have any filters assigned to it.

Factory Device Collections

Data Aggregator and CA Performance Center support the concept of *device collections*, which are logical groupings of monitored devices.

Several factory (out-of-the-box) device collections are provided to get data into your Data Aggregator system quickly and test the product. Devices that are detected during discovery are added to these device collections depending on their type. For example, routers are added to the factory All Routers device collection. Upon synchronization, these monitored devices are added to the corresponding device collections in CA Performance Center.

Factory monitoring profiles are then automatically applied to factory device collections, allowing data to be collected immediately without any intervention on your part. Once this data has been collected, you can run reports on the data to gain a better understanding of your network.

The following factory device collections are provided:

- [All Devices](#) (see page 87)
- [All Routers](#) (see page 87)
- [All Servers](#) (see page 87)
- [All Switches](#) (see page 88)
- [All Manageable Devices](#) (see page 88)
- [All ESX Hosts](#) (see page 88)
- [All Virtual Machines](#) (see page 89)
- [All VMware vCenters](#) (see page 89)

Note: The factory device collections are mostly for use in a lab or in a demo setting. In a real production deployment, the best practice is to design and configure custom device collections to have granular control and optimal data collection.

Access the Monitoring Configuration menu to see a list of device collections and to see the monitoring profiles that are applied to each. Administrators can view the device collections for the tenant they are administering. A tenant administrator can view its own list of device collections.

More information:

[Factory Monitoring Profile Associations](#) (see page 81)

[Troubleshooting: Polling Has Stopped on Discovered Metric Family](#) (see page 154)

[Custom Device Collections](#) (see page 89)

All Devices Device Collection

The All Devices device collection is a factory device collection. Manageable and pingable devices that are detected during a discovery are automatically placed into the All Devices device collection. Inaccessible devices are not included in the All Devices device collection.

Important! Do not associate monitoring profiles with the All Devices device collection. Doing so can cause extra SNMP requests being made to pingable-only devices, and can result in sporadic metric family support.

More information:

[Factory Device Collections](#) (see page 85)

[Factory Monitoring Profile Associations](#) (see page 81)

All Routers Device Collection

The All Routers device collection is a factory device collection. Routers that are detected during a discovery are automatically placed into the All Routers device collection.

Note: Routers can appear in both the All Routers device collection and the All Switches device collection.

More information:

[Factory Device Collections](#) (see page 85)

[Factory Monitoring Profile Associations](#) (see page 81)

All Servers Device Collection

The All Servers device collection is a factory device collection. Physical and virtual servers (hosts) that are detected during a discovery are automatically placed into the All Servers device collection. Network devices such as routers and switches are not included in the All Servers device collection.

More information:

[Factory Device Collections](#) (see page 85)

[Factory Monitoring Profile Associations](#) (see page 81)

All Servers Device Collection

The All Switches device collection is a factory device collection. Switches that are detected during a discovery are automatically placed into the All Switches device collection.

Note: Switches can appear in both the All Routers device collection and the All Switches device collection.

More information:

[Factory Device Collections](#) (see page 85)

[Factory Monitoring Profile Associations](#) (see page 81)

All Manageable Devices Device Collection

The All Manageable Devices device collection is a factory collection. Manageable devices collect advanced performance statistics and are monitored with a protocol such as SNMP. Manageable devices that are detected during a discovery are automatically placed into the All Manageable Devices device collection.

Pingable devices can only be monitored for availability and do not provide any additional performance metrics. Therefore, pingable devices are not included in the All Manageable Devices device collection.

Note: Manageable devices can appear in both the All Devices device collection and the All Manageable Devices device collection.

More information:

[Factory Device Collections](#) (see page 85)

All ESX Hosts Device Collection

The All ESX Hosts device collection is a factory device collection. ESX hosts that are detected during discovery are placed into the All ESX Hosts device collection automatically.

More information:

[Factory Device Collections](#) (see page 85)

[Factory Monitoring Profile Associations](#) (see page 81)

All Virtual Machines Device Collection

The All Virtual Machines device collection is a factory device collection. VMware virtual machines that are detected during discovery are placed into the All Virtual Machines device collection automatically.

More information:

[Factory Device Collections](#) (see page 85)

[Factory Monitoring Profile Associations](#) (see page 81)

All VMware vCenters Device Collection

The All VMware vCenters device collection is a factory device collection. All servers that are running systemEdge with the VCAIM that are detected during discovery are placed into the All VMware vCenters device collection automatically.

More information:

[Factory Device Collections](#) (see page 85)

[Factory Monitoring Profile Associations](#) (see page 81)

Custom Device Collections

The factory device collections are mostly for use in a lab or in a demo setting. In a real production deployment, the best practice is to design and configure custom device collections to have granular control over what is being polled. For example, disable polling of a device by disassociating the device from any other device collection that has monitoring profiles that are associated to it. If you are associating monitoring profiles to the factory device collections (such as All Routers), then you cannot stop a single device from being polled. Devices cannot be removed from factory device collections, so you disassociate the monitoring profiles instead to disable polling. You then create custom device collections that contain devices that you want to apply the same polling policy to. Associate monitoring profiles (or custom monitoring profiles) to those custom device collections to begin polling.

Create custom device collections in CA Performance Center, then either synchronize them immediately with Data Aggregator, or wait for the automatic synchronization. Upon synchronization, Data Aggregator creates the corresponding device collections for use in monitoring devices.

Note: For more information about creating custom device collections and synchronizing them with Data Aggregator, see the *CA Performance Center Administrator Guide*.

Access the Monitoring Configuration menu to see a list of device collections and to see the monitoring profiles that are applied to each. Administrators can view the device collections for the tenant they are administering. A tenant administrator can view its own list of device collections.

More Information:

[Factory Device Collections](#) (see page 85)

View Monitored Devices

You can view details for monitored devices and can view their associations with device collections, components, monitoring profiles, and metrics. You can also view a Filter Report. This information helps you to see information in context, such as which monitoring profiles are being used to poll device components.

Note: Some features require administrator privileges.

The monitored devices are manageable or pingable (accessible but not manageable). The inaccessible devices are not monitored devices. Components of monitored devices can be viewed from the Polled Metric Families tab.

Follow these steps:

1. Click Monitored Devices from the Monitored Inventory menu for a Data Aggregator data source.

The Tree View tab displays.

2. Select Device by Collection or Device by Monitoring Profile from the drop-down list, and select a specific device from the corresponding tree view.

Note: Alternatively, select the Search tab to search by host name, device name, or IP address. You can enter a partial name or IP address to return a list of devices that contain that partial match. Wildcards and regular expressions are not supported.

The Details tab provides details for the selected monitored device. Details include the device IP address, associated SNMP profile, the status of the device, and so on. You can edit the IP address, Data Collector host, SNMP profile, and SNMP version for the device.

You can edit the device IP address in two ways:

- Edit the IP Address field and then click Save.
- Right-click an IP address in the IP Addresses table, and select 'Set this IP as the primary IP for the device.' Click Save.

Note: More information is available in this view for manageable devices.

(Optional) Click Rediscover to rediscover the device. The following set of attributes can be updated as a result of this discovery:

- System name
- Hostname
- Device type (as it appears in CA Performance Center)
- Location
- Vendor
- Device description
- Device model

Note: Changes to device attributes can result in changes to the groups and device collections a device is a part of. Changes to groups and device collections can potentially add or remove monitoring profiles.

Verify that the device was rediscovered by looking for the event that the rediscovery triggered. To view events, click the Dashboards menu in CA Performance Center and select "Events Display" under Operations Displays.

3. (Optional) Select a metric family and click Update Metric Family to reconfigure components for any configuration updates. For example, if you add a disk drive on a server, you can use the Update Metric Family button to rediscover the configuration update. The configuration update creates a disk component.
4. Select another tab:

- The Polled Metric Families tab shows the total set of metric families that are polled on a device and shows their poll rates. This total set is based on the consolidation of all of the monitoring profiles on the device. The tab also shows whether the device supports the metric family.

The Components table for a given metric family shows the polling status on the components for a metric family component that was previously discovered. One of the following displays in the Status column:

Active

Indicates that the component is being polled.

Inactive

Indicates that polling has stopped on the component because the metric family is no longer monitored for the device.

Retired

Indicates that the component no longer exists on the physical device. Polling is stopped on the component. You can view historical data for reporting purposes. By default, retired components are not synchronized with CA Performance Center. To enable this option, select the Synchronize retired items checkbox on the Edit Data Source dialog on the Manage Data Sources page in CA Performance Center.

(Optional) Select a metric family and click Update Metric Family to reconfigure components for any configuration updates. For example, if you add a disk drive on a server, you can use the Update Metric Family button to rediscover the configuration update. The configuration update creates a disk component.

- The Threshold Profiles tab shows the threshold profiles that are applied to the selected device due to the groups to which the device belongs.
- The Monitoring Profiles tab lets you select a device collection to see the associated profiles names. Hover over a profile to see the description.
- The Metrics tab is populated with a list of metrics that this device supports. Select a metric family to view its details. View the backing vendor certification, vendor source (the MIB table source is displayed if it is an SNMP vendor certification), and the expression that is used to calculate each metric.
- The Filter Report tab shows which interface filter criteria have been used during component monitoring. The tab also shows a report of all of the interfaces that are identified on the device and whether they matched the specified filter criteria. If you change the rules on a custom monitoring profile, the Interface Filter Criteria pane does not reflect those changes. If you disassociate the monitoring profile from a group, the Interface Filter Criteria pane does not reflect those changes. Rediscover the device to filter the interfaces that are based on the changes you made to the filter criteria and monitoring profile.

More Information:

[Discovery and Polling](#) (see page 57)

[Update Device Reconfiguration Manually](#) (see page 103)

[Clear an Interface Filter](#) (see page 116)

[Troubleshooting: Polling Has Stopped on Discovered Metric Family](#) (see page 154)

[How to Set and Activate an Interface Filter](#) (see page 115)

Delete Devices

You can delete discovered devices. For example, you can delete a discovered device if you want to stop monitoring it.

When you delete a device, the following results occur:

- All of the associated device components are deleted.
- Historical data on the deleted devices and device components is no longer accessible.

Note: If any existing discovery profiles are rerun, the deleted devices can be discovered again.

Follow these steps:

1. Click Monitored Devices from the Monitored Inventory menu for a Data Aggregator data source.

The Tree View tab displays.

2. Select the Search tab.

Note: Do not use the global Search box at the top of the page.

3. To search for the monitored device you want to delete, enter text in the local Search box. You can search by host name, by device name, or by IP address. You can enter a partial name or IP address to return a list of devices that contain that partial match.

Note: Wildcards and regular expressions are not supported.

A list of matching devices is returned.

4. Do one of the following steps:

- Select the single monitored device or multiple devices that you want to delete and click Delete.
- Delete *all* of the devices in the results list by selecting the checkbox that appears next to the Name column and clicking Delete.

A confirmation dialog opens.

5. Click Yes to confirm the deletion.

The devices are deleted and are no longer visible in the Monitored Devices inventory. If another data source is not managing these devices, the devices no longer appear in the Inventory view or as members of groups the next time Data Aggregator synchronizes with CA Performance Center.

Change the Primary IP Address for a Monitored Device

You can change the primary IP address for a monitored device.

Follow these steps:

1. Click Monitored Devices from the Monitored Inventory menu for a Data Aggregator data source.

The Tree View tab displays.

2. Select Device by Collection or Device by Monitoring Profile from the drop-down list, and select a specific device from the corresponding tree view.

Note: Alternatively, select the Search tab to search by host name, device name, or IP address. You can enter a partial name or IP address to return a list of devices that contain that partial match. Wildcards and regular expressions are not supported.

3. Change the primary IP address by taking one of the following steps:
 - Edit the IP Address field and then click Save.
 - Right-click an IP address in the IP Addresses table, select 'Set this IP as the device's primary IP,' and then click Save.

The primary IP address is changed.

Delete Retired Components

Data Aggregator includes a script to delete retired components. Retired components are components that no longer exist on physical devices. The presence of excessive numbers of retired components can impact user interface performance. To delete retired components, understand how to use this script.

Note: To automate the removal of retired components, see the *Data Aggregator Administration Using REST Web Services Guide*.

Follow these steps:

1. Open a command prompt and access the /opt/IMDataAggregator/scripts directory.
2. To call the script to delete retired components, type the following command:

```
./remove_retired_items.sh
```

The script parameters are listed and are described.

Example: Return a Total Number of Retired Components

1. Type the following command:

```
./remove_retired_items.sh -h host_name
```

-h *host_name*

Specifies the Data Aggregator host name to connect to.

The total number of retired components is displayed.

2. (Optional) Enter the number 1 to return a list of the *names* of the retired components.
3. (Optional) Enter the number 1 to delete all of the retired components.

Example: Filter the List of Retired Components By a Specific Criterion

1. To call the script to delete retired components, type the following command:

```
./remove_retired_items.sh
```

The script parameters are listed and are described.

2. Delete retired components by a specific criterion:

- If multiple Data Collector instances are installed, you can have duplicate IP addresses. To filter retired components by IP address, do the following steps:

a. Type the following command:

```
./remove_retired_items.sh -h host_name -a device_IP_address
```

b. **Note:** You cannot enter a range of IP addresses.

c. (Optional) Enter the number 1 to return a list of the *names* of the retired components.

d. (Optional) Enter the number 1 to delete all of the retired components.

- To delete by the age of the retired components (in number of days old), type the following command:

```
./remove_retired_items.sh -h host_name -t  
filter_by_days_old_from_current_time
```

For example, the following command deletes retired components that are less than ten days old from the current time:

```
./remove_retired_items.sh -h host_name -t 10
```

Example: Delete a Large Number of Retired Components

You can easily delete all of the retired components where the total exceeds 100,000:

- To review all of the retired components where the total exceeds 100,000, type the following command:

```
./remove_retired_items.sh -h host_name -o outputfile
```

-o *outputfile*

Is the output of all of the retired components. The output is a .csv file.

For example, the following command outputs a list of all of the retired components. The .csv file format includes the device item ID, the device display name, the retired component ID, and the retired component display name:

```
./remove_retired_items.sh -h my_host_name -o myretired.csv
```

- To delete all of the retired components where the total exceeds 100,000, and to log the information to a .csv file, type the following command:

```
./remove_retired_items.sh -h host_name -o outputfile -c Yes
```

-o *outputfile*

Is the output of all of the retired components. The output is a .csv file.

-c Yes

Confirms the deletion of all retired components.

For example, the following command deletes all of the retired components:

```
./remove_retired_items.sh -h my_host_name -o myretired.csv -c Yes
```

Consider further details about deleting retired components:

- If you filter retired components by IP domain name or IP domain ID, also specify a specific IP address to return correct results.
- If your filter criteria return too many retired components, the REST interface does not return a response. Use other filtering options to narrow the results. More filter criteria are available at <http://hostname:port/rest/retired/xsd/filterselect.xsd>.

IP Domain Deletions

You can delete IP domains. For example, you can delete an IP domain when merging two or more domains. You can also delete an IP domain that was used for testing purposes. Deleting an IP domain deletes all devices and device components that are associated with it. Deleting an IP domain also invalidates the discovery profile that is associated with that IP domain.

You delete IP domains in CA Performance Center. After you delete an IP domain, synchronize the deletion with Data Aggregator, or wait for the automatic synchronization to occur.

Note: For more information about deleting and synchronizing IP domains, see the *CA Performance Center Administrator Guide*.

Once Data Aggregator is aware that an IP domain has been deleted, the following results occur:

- All devices and device components that are associated with the deleted IP domain are deleted.
- Data Collector that is associated with the deleted IP domain is stopped. The status reads 'Not Collecting Data'.

Note: You can delete an IP domain when Data Collector is down. When Data Collector comes back up, an error message displays in the *Data Collector installation directory/apache-karaf-2.3.0/shutdown.log* file and Data Collector shuts down immediately.

- All discovery profiles that specify the deleted IP domain are invalidated and cannot be run. The state reads 'No IP Domain'. All running Discoveries in a deleted IP domain are aborted.

Note: You can change the state of an invalidated discovery profile back to the 'Ready' state by specifying a valid IP domain for which you want to discover devices.

- An audit event is generated on the associated tenant item for each deleted device.

More information:

[Discovery Profiles](#) (see page 60)

Tenant Deletions

You can delete tenants. For example, you can delete a tenant if you are a managed service provider (MSP) and a tenant is no longer your customer. Deleting a tenant deletes all devices, device components, IP domains, SNMP profiles, and discovery profiles that you associated with it.

Note: You cannot delete the Default Tenant.

You delete tenants in CA Performance Center. After you delete a tenant, manually synchronize the deletion with Data Aggregator, or wait for the automatic synchronization to occur.

Note: For more information about deleting and synchronizing tenants, see the *CA Performance Center Administrator Guide*.

Once Data Aggregator is aware that a tenant has been deleted, the following events occur:

- All devices, device components, IP domains, SNMP profiles, and discovery profiles that are associated with the deleted tenant are deleted.
- Polling on the deleted devices and device components is stopped.
- Historical data on the deleted devices and device components is no longer accessible.
- An audit event is generated on the Data Aggregator device for each deleted tenant.
- All threshold events on the deleted device and its deleted components are removed.

Note: You can delete a tenant when Data Collector is down. An error message displays in the *Data Collector installation directory/apache-karaf-2.3.0/shutdown.log* file when Data Collector comes back up and then Data Collector shuts down immediately.

Disable Tenants

You can disable tenants. For example, you can disable a tenant if you are a managed service provider (MSP) and you want to stop active monitoring of the tenant infrastructure.

Note: Log in as the administrator to perform this task.

You disable tenants in CA Performance Center. After you disable a tenant, synchronize the disablement with Data Aggregator, or wait for the automatic synchronization to occur.

Note: For more information about disabling tenants, see the *CA Performance Center Administrator Guide*.

Once Data Aggregator is aware that a tenant has been disabled, the following results occur:

- The Data Aggregator system stops all Data Collector hosts that are associated with the disabled tenant. Data Collector hosts then show a status of 'Not Collecting Data'. (When the tenant is reenabled, the Data Collector hosts must be restarted manually.)

Note: For any new Data Collector installations for a disabled tenant, the status 'Not Collecting Data' is shown. Discovery is only permitted when the tenant is enabled again.

- All devices, device components, IP domains, SNMP profiles, and discovery profiles that are associated with the disabled tenant continue to exist.
- Polling is stopped for any devices and components that are being monitored on behalf of the disabled tenant.
- Historical data on the devices and components for a tenant remains accessible.
- Discovery profiles that are associated with the disabled tenant are invalidated and cannot be run. The discovery profiles have a state of 'Tenant Disabled'.
- If a discovery profile is invalidated while a discovery is running on it, the discovery is aborted.
- An audit event is generated on the Data Aggregator device for the disabled tenant.

More Information:

[Enable Tenants](#) (see page 99)

Enable Tenants

You can enable a tenant that you previously disabled. For example, you can enable a tenant if you are a managed service provider (MSP) and you want to restart active monitoring of the tenant infrastructure.

Note: Log in as the administrator to perform this task.

You enable tenants in CA Performance Center. After you enable a tenant, do the following actions:

1. Synchronize the activation with Data Aggregator, or wait for the automatic synchronization to occur.

Note: For more information about enabling tenants, see the *CA Performance Center Administrator Guide*.

The following results occur:

- Data Aggregator becomes aware that a tenant has been enabled.
 - Discovery profiles that are associated with the enabled tenant are validated. The discovery profiles display their current state.
 - An audit event is generated on the Data Aggregator device for the tenant.
2. [Manually restart all Data Collector hosts that are associated with the tenant](#) (see page 48).

The following results occur:

- Polling is restarted for any devices and components that are being monitored on behalf of the enabled tenant.
- Discovery profiles that are associated with the enabled tenant can be run.

More information:

[Disable Tenants](#) (see page 98)

[Troubleshooting: Discovery Does Not Start](#) (see page 153)

Device Reconfiguration

Device reconfiguration changes can be monitored and updated automatically or manually in Data Aggregator to keep device components up-to-date. Device reconfiguration includes changes to physical device components and software configuration changes, such as monitoring response path tests for protocols. Data Aggregator uses the same method to monitor both types of reconfiguration.

Additional examples of reconfiguration changes include:

- Adding a board to a device, which adds more ports to the device.
- Adding memory, CPUs, physical interfaces, or any metric family to a discovered device.
- Reconfiguring a virtual switch.
- Changing the configuration of a device so that a discovered device participates in routing protocols.

When change is detected, Data Aggregator generates reconfiguration events and can update its representation of the metric family to reflect the changes to device components. View reconfiguration events by selecting Dashboards, Operations, Events Display.

Understanding how change detection works in Data Aggregator helps you to select the options that are best suited to monitoring device reconfiguration in your environment. For example, you can set the frequency for change detection monitoring.

More Information:

[How to Manage Change Detection](#) (see page 100)

How to Manage Change Detection

Change detection management planning helps ensure that Data Aggregator detects and monitors device reconfigurations in your environment according to your needs. You can plan ahead for any device reconfiguration when you first set up Data Aggregator to discover new devices. Or, you can edit these options at any time after devices are discovered.

The choices that you make are based on:

- The likelihood of change.
- The frequency of change you anticipate.
- Your tolerance for how far out-of-date your data is.

There can be metric families, such as CPUs, that you want to monitor for reconfiguration infrequently. For other metric families that are more dynamic, such as virtual systems, choose a more frequent rate.

The basic process for setting change detection is:

1. Create or edit a *custom* monitoring profile. (You can also copy a factory monitoring profile and edit the copy.)
2. Select Enable Change Detection and set the Change Detection Settings, Rate in the monitoring profile.

The Change Detection Settings, Rate option is used to set the frequency at which Data Aggregator checks for changes. The rate of detection can be set in minutes or hours. By default, the rate is set to 24 hours.

Note: Consider how frequently the metric family is likely to change, and how many devices the monitoring profile is applied to. Avoid setting change detection rates that are more frequent than necessary.

3. Update the Data Aggregator representation of the metric families.

After you set the change detection rate, you have two options for correcting the Data Aggregator configuration: automatic or manual update of the metric families. This option does not update the metric families. Instead, it updates the representation of the metric families by making sure that the correct set of components is being monitored.

- Selecting the Automatically Update Metric Families option (selected by default) means that you do not have to intervene when a reconfiguration is detected. Data Aggregator automatically starts monitoring any new components and retires any components that are no longer detected when a reconfiguration event occurs.

View the Events Display dashboard to see reconfiguration events:

- If the components have changed for a device, an event is generated on the associated device. This event describes that a component change has been detected and will be applied after a short period.
- After component reconciliation is applied, another event is generated. This event describes how many components were added, retired, and remained unchanged.

- Deselecting the Automatically Update Metric Families option means that Data Aggregator does not automatically start monitoring new components or retire old components.

View the Events Display dashboard to see reconfiguration events:

- If the components have changed for a device, an event is generated on the associated device. This event describes that a component change has occurred but reconciliation has not occurred.

To have the reconfiguration changes applied, manually click the Update Metric Family button on the Polled Metric Families page for a device.

4. To make the monitoring profile active, assign the custom monitoring profile to a device collection.

Examples:

- If you are aware that your environment will be undergoing significant maintenance, you can turn off the automatic update until the major maintenance is complete. For small, regular changes, enabling the automatic update feature helps ensure your Data Aggregator stays up to date.
- Monitoring profiles are assigned to device collections that contain devices. If you have special devices that you want to monitor differently, create a custom device collection for these special devices and assign a custom monitoring profile with the change detection settings you want. For example, you can monitor critical core routers more frequently than other routers by creating a critical core routers device collection and assigning a custom monitoring profile that performs change detection hourly. Other routers can remain in the 'All Routers' device collection using the factory monitoring profile (with no change detection) or a custom monitoring profile that you set to less frequent change detection.

More information:

[Discovery and Polling](#) (see page 57)

[Update Device Reconfiguration Manually](#) (see page 103)

[Update Device Reconfiguration Automatically](#) (see page 102)

Update Device Reconfiguration Automatically

Reconfiguration changes to a discovered device can affect the metric families that are associated with the device. The device reconfiguration can be set to update automatically in the monitoring profile the metric families are assigned to, which applies to any metric families that the monitoring profile includes. This option is set by default when you create a custom monitoring profile, but it can also be edited at any time. This procedure tells you how to set the Automatically Update Metric Families option in an existing custom monitoring profile if it has been previously deselected.

When the metric family is updated, Data Aggregator has an accurate representation of the device configuration. Reports that you generate reflect accurate information.

Follow these steps:

1. Navigate to the list of all monitoring profiles.
2. Select the monitoring profile that you want to update automatically and click Edit.

3. Select Enable Change Detection and then do the following steps:

- Set the Change Detection Settings, Rate to a value that is greater than zero.

Note: Consider how frequently the metric family is likely to change, and how many devices the monitoring profile is applied to. Avoid setting change detection rates that are more frequent than necessary.

- Select Automatically Update Metric Families.
- Click Save.

When you make a configuration change to a device associated with this monitoring profile, the device configuration is updated automatically.

When a device configuration is updated, Data Aggregator does the following steps:

- Generates an event on the monitored device.
- Identifies new components and creates them.
- Identifies components that are no longer present and retires them.

Note: By default, retired components are not synchronized with CA Performance Center. To enable this option, select the Synchronize retired items checkbox on the Edit Data Source dialog on the Manage Data Sources page in CA Performance Center.

- Identifies existing components that have changed from a previous discovery. The Name column changes, if applicable.

Note: Historical data is accessible and can be reported on.

More information:

[How to Manage Change Detection](#) (see page 100)

Update Device Reconfiguration Manually

Reconfiguration changes to a discovered device can affect the metric families that are associated with the device. The device reconfiguration can be updated manually when the Automatically Update Metric Families option is not selected in the associated monitoring profile. In this case, you view the event logs to identify reconfiguration events that you want to update the metric families for.

When the metric family is updated, Data Aggregator has an accurate representation of the device configuration. Reports that you generate reflect accurate information.

Follow these steps:

1. [View the event logs to identify reconfiguration events for which you want to update the metric families](#) (see page 132).

2. Click Monitored Devices from the Monitored Inventory menu for a Data Aggregator data source.

The Tree View tab opens.

3. Select Device by Collection from the drop-down list, and select the monitored device that was updated from the corresponding tree view.

The Polled Metric Families tab shows the consolidated monitoring profiles that are associated with a device. Devices only have one consolidated monitoring profile. Each consolidated monitoring profile lists every metric family that can be polled on the device and whether the device supports the metric family.

4. Select the metric family for which you want to update the configuration and click Update Metric Family.

Your device configuration is updated and Data Aggregator does the following steps:

- Generates an event on the monitored device.
- Identifies new components and creates them.
- Identifies components that are no longer present and retires them.

Note: By default, retired components are not synchronized with CA Performance Center. To enable this option, select the Synchronize retired items checkbox on the Edit Data Source dialog on the Manage Data Sources page in CA Performance Center.

- Identifies existing components that have changed from a previous Discovery. The Name column changes, if applicable.

Note: Historical data is accessible and can be reported on.

More information:

[View Monitored Devices](#) (see page 90)

[How to Manage Change Detection](#) (see page 100)

Chapter 5: Managing Interfaces

This section contains the following topics:

[How to Poll Critical Interfaces Faster than Noncritical Interfaces](#) (see page 105)

[How to Set and Activate an Interface Filter](#) (see page 115)

[Clear an Interface Filter](#) (see page 116)

[Interface Components Naming Convention](#) (see page 117)

[Interface Utilization Calculation](#) (see page 117)


How to Poll Critical Interfaces Faster than Noncritical Interfaces

As the Administrator, you need frequent data about your most critical systems while maximizing the overall performance of your performance management systems. One way to accomplish your goal is by polling only critical interfaces at a high rate, while polling noncritical interfaces at a normal or slow rate. You can poll at differing rates by using a filter on the Interfaces metric family that is associated with your monitoring profile. By fast-polling interfaces sparingly, you can reduce unnecessary network traffic and performance management system load while still sufficiently monitoring the health of your network system.

For example, your data center access switch connects many application servers to only two aggregation switches. You decide to poll the interfaces supporting these aggregation switches at a higher rate. These links are critical, because they support network traffic to all other connected switches. However, polling *all* interfaces at a higher rate would cause unnecessary network traffic, wasting system resources and possibly causing network performance issues. After consulting with your network operations and engineering teams, you decide that a normal polling rate is sufficient for the interfaces connecting each attached server. To apply different polling rates, you implement two monitoring profiles for interfaces.

Note: Filters that you set on metric families are ignored when event rules that are applied to monitoring profiles trigger events.

The following illustration shows how to configure monitoring profiles to poll interfaces at varying rates:



Procedures

[View the existing monitoring profiles](#) (see page 107).

[Copy the factory Network Interface monitoring profile](#) (see page 108).

[Set a filter on the Interface metric family](#) (see page 110).

Procedures

[Associate the monitoring profile with a device collection](#) (see page 112).

[View the monitored devices to verify your results](#) (see page 113).

Note: For more information about how monitoring profiles work with device collections and metric families, see the *Data Aggregator Overview Guide*.

View Your Monitoring Profiles

As the CA Performance Center Administrator, you decide to poll critical interfaces as often as possible. However, you want to minimize unnecessary network traffic that polling *all* interfaces at this fast rate can produce. You decide to create two monitoring profiles for interfaces—one with normal polling, and one with fast polling.

Before you create a monitoring profile, you review the existing monitoring profiles to find one that closely matches your needs.

Follow these steps:

1. Click Monitoring Profiles from the Monitoring Configuration menu for your Data Aggregator data source.

A list of monitoring profiles is populated.

2. Select a monitoring profile.

Details for the selected monitoring profile populate the tabs:

- Metric Families tab—Displays a list of metric families that are associated with that specific monitoring profile. Metric families contain the metrics that are used for polling devices and components.
- Collections tab—Displays a list of device collections that are associated with that specific monitoring profile.

More information:

[Troubleshooting: Polling Has Stopped on Discovered Metric Family](#) (see page 154)

Copy a Factory Monitoring Profile

As the CA Performance Center Administrator, you find that the factory Network Interfaces monitoring profile closely matches your needs and requires only minor changes. Therefore, you create a copy and use it to poll only critical interfaces at a faster polling rate.

Note: Log in as the administrator to perform this task.

Follow these steps:

1. [Navigate to the list of all monitoring profiles in CA Performance Center](#) (see page 107).
2. Select the Network Interfaces monitoring profile and click Copy.

Note: Factory monitoring profiles cannot be edited or deleted. All monitoring profiles, including custom, are global.

The Create/Edit Monitoring Profile dialog opens.

3. Enter the following information for your monitoring profile:

- **Name:** Uplink Interfaces
- **Description** (optional): Monitors performance of interfaces in all critical Uplink devices.
- **SNMP Poll Rate:** 1 minute

Note: We recommend that you rename the profile. Unique naming is enforced across all tenants.

Consider the following information about poll rates:

- When the poll rate is changed, it takes up to two cycles for the new poll rate to take effect. When the 60-minute rate is used to poll an existing device, a 'No Data To Display' message appears in the dashboard view given the default time range of Last Hour. If you change the dashboard setting to a prior hour, it is possible to see earlier data. However, the view does not display the latest data until the new poll cycle completes.
- Interfaces that are assigned to multiple monitoring profiles with different poll rates are polled at the fastest assigned rate.

4. Leave the Change Detection Settings, Rate value at 24 Hours.

Consider the following information about change detection rates:

- The *change detection rate* is how often Data Aggregator checks whether any components on a device have been reconfigured. Changes can include new components that have been created or existing components that have been retired.

Note: The reconciliation algorithm specified in the metric family defines the configuration changes to watch for. For more information about how change detection and device reconfiguration works, see the *Data Aggregator Administrator Guide*.

- The Change Detection Settings, Rate option is used to set the frequency at which Data Aggregator checks for changes. The rate of detection can be set in minutes or hours. By default, the rate is set to 24 hours.
- Changes are detected at the fastest rate you specified for all of the monitoring profiles that are associated with a collection of devices.

5. Leave the 'Automatically Update Metric Families' check box selected.

This option controls the Data Aggregator response once a change or reconfiguration is detected. Selecting this option automatically causes Data Aggregator to start monitoring new components or to stop monitoring retired components. When this option is not selected, you can manually control monitoring of components, as follows:

- a. Manually check the Events Display dashboard to watch for configuration events.
- b. Navigate to the Data Aggregator administration menu, Monitored Devices, Polled Metric Families view.
- c. Select the appropriate metric family, and click Update Metric Family to help ensure that Data Aggregator picks up the latest device reconfiguration.

Note: If an interface filter is applied, Data Aggregator monitors only the interfaces that pass the filter conditions after reconfiguration.

6. Leave the Interfaces metric family as the only metric family in the Selected Metric Families list.
7. Click Save.

Your copied monitoring profile is added to the Monitoring Profiles list. However, this monitoring profile is not active until you assign it to a device collection.

More information:

[Discovery and Polling](#) (see page 57)

[View Events](#) (see page 132)

[How to Manage Change Detection](#) (see page 100)

[Troubleshooting: Polling Has Stopped on Discovered Metric Family](#) (see page 154)

[Discovery Workflow](#) (see page 54)

[How to Set and Activate an Interface Filter](#) (see page 115)

Set an Interface Filter

By default, the factory network interface monitoring profile includes a filter to prevent modeling interfaces that are administratively down. In addition, interfaces with a type (ifType) of 1 (Other) or 24 (Loopback) are not modeled, regardless if those interfaces are administratively up or down. IPSLAs with the rttMonCtrlAdminOwner MIB object that contains the string "Network Health" are not modeled either.

Filtering reduces the number of interfaces that are monitored, which reduces unwanted data collection and network traffic.

In addition to polling only administratively up interfaces, you also want to poll the most critical interfaces more frequently. To isolate and poll only these interfaces faster, you add a second filter condition to the interface filter associated with your custom monitoring profile. This second filter condition isolates the critical interfaces by finding only interfaces that contain "uplink" in their description.

Note: Log in as the administrator to perform this task.

Follow these steps:

1. [Select your interfaces monitoring profile \(called "Uplink Interfaces"\) from the Monitoring Profiles page](#) (see page 108).

2. Click Interface metric family row on the Metric Families tab and click Edit Filter.

Note: Do not click directly on the metric family name, because it is linked to take you to the metric family definition. Instead, click the row the metric family name is in to activate the Edit Filter option.

3. Click the Add Condition button.

Note: Multiple conditions are connected with an "and" operation. That is, all conditions must be met to satisfy the filter.

4. Configure the filter conditions with the following options and click Save:

- Attribute: Description
- Operation: Contains
- Filter Value: uplink

Note: The Filter Value field is case-sensitive.

Consider the following details about additional attributes you can use for filtering:

- For Speed In and Speed Out, you can use a decimal in the text field (such as 1.544) and can specify bps, Kbps, Mbps, or Gbps.
- For more information about configuring Type (that is, ifType), see the iana web site: <http://www.iana.org/assignments/ianaiftype-mib>
<http://www.iana.org/assignments/ianaiftype-mib>.
- For Description and Alias, you can use a regular expression for filtering only when you select the Matches Regex or the Does Not Match Regex operation.

When you save your changes, the filter criteria display on the Metric Families tab. You can now apply this monitoring profile to the appropriate device collection to begin polling your selected interfaces.

Note: Data Aggregator applies filtering after discovery. Interface components that do not match the filter criteria are not polled. If you add or edit an Interface filter *after* you run a discovery, polling on these components stops. These interface components are *not* displayed in CA Performance Center dashboards and data views.

Considerations for Interface Filters and Multiple Monitoring Profiles

When multiple monitoring profiles are assigned to a device collection, the filter matching criteria follows the "or" rule. So, Data Aggregator monitors all interfaces that satisfy the criteria for any of the monitoring profiles in the group.

Some of the monitoring profiles may have filters and some may not. Plus, these profiles can specify differing poll rates. In this case, Data Aggregator monitors the interfaces that match any monitoring profile, but the polling rates can differ. If more than one monitoring profile applies to an interface, Data Aggregator polls the interface once, and polls it at the fastest polling rate:

- Monitoring Profile 1—Filter: Description contains "X," Poll Rate: 1 minute
- Monitoring Profile 2—Filter: None, Poll Rate: 5 minutes
- Monitoring Profile 3—Filter: Description contains "Y," Poll Rate: 10 minutes

In this example, interfaces that match Monitoring Profile 1 are polled every minute. All other interfaces are polled every 5 minutes. Interfaces that match Monitoring Profile 3 also match Monitoring Profile 2, which does not include a filter. The fastest poll rate applies, so no interfaces are polled at 10-minute intervals.

In this case, if one monitoring profile has no filter, the result is that many interfaces may be polled more frequently than necessary. Therefore, after you set a filter, remove associations from other monitoring profiles to make sure that only components matching the specified filter are monitored.

Assign Your Monitoring Profile to a Device Collection

As the administrator or a tenant administrator, you associate the new Uplink Interfaces monitoring profile with a device collection to begin polling. In this case, you associate the profile with the Switches device collection, which is the same device collection that is associated with the factory Network Interfaces monitoring profile. Polling rates are applied to the interfaces in this device collection, as follows:

- Fast polling: Interfaces that satisfy the filter criteria of the Uplink Interfaces monitoring profile.
- Normal polling: All other interfaces that the Network Interfaces monitoring profile discovers.

Important! All custom monitoring profiles are global and visible to tenant administrators. However, the association of a monitoring profile with a specific device collection can be scoped to a tenant.

Follow these steps:

1. Click Collections from the Monitoring Configuration menu for your Data Aggregator data source.

A list of device collections displays. Administrators can view the device collections for the tenant they are administering. A tenant administrator can view its own (tenant) list of device collections.

2. Select the All Switches device collection and click the Monitoring Profiles tab.

A list displays the monitoring profiles that are associated with the selected device collection. The Network Interface device collection exists in this list.

3. Click Manage.

The Assign Collection Monitoring Profiles dialog opens.

4. Select the Uplink Interfaces monitoring profile and click Add.

The selected monitoring profile moves to the Assigned Monitoring Profiles list.

5. Click Save.

Your changes are saved.

View Monitored Devices to Verify Results

After you set up your monitoring profiles, review the monitored devices and the Filter report to verify that only your critical devices are polled at the higher rate. This information helps you to see information in context, such as which monitoring profiles are being used to poll device components. Verifying the results can help you identify any necessary adjustments to help you achieve the polling results that you want.

Note: Monitored devices are manageable devices and pingable (accessible but not manageable). Inaccessible devices are not monitored devices. Components of monitored devices can be viewed from the Polled Metric Families tab.

Follow these steps:

1. Run an on-demand discovery.

Note: If your discovery profile runs automatically, you can wait for the next scheduled discovery. For more information about managing discovery, see the *Data Aggregator Administrator Guide*.

2. Click Monitored Devices from the Monitored Inventory menu for a Data Aggregator data source.
3. Select one of these options from the drop-down list to locate one of your aggregation switch devices in the corresponding tree view:
 - Device by Collection—Your devices appear under the All Switches device collection.
 - Device by Monitoring Profile—Your critical interfaces appear under Devices under the Uplink Interfaces monitoring profile.

Note: Alternatively, select the Search tab to search by host name, device name, or IP address. You can enter a partial name or IP address to return a list of devices that contain that partial match. Wildcards and regular expressions are not supported.

The Polled Metric Families tab shows the consolidated monitoring profiles that are associated with the switch device. Devices have only one consolidated monitoring profile. Each consolidated monitoring profile lists every metric family to poll on the device and whether the device supports the metric family.

4. Select the Interface metric family.

The Components table for the Interfaces metric family shows one of the following polling statuses for the discovered Interface components:

Active

Indicates that the component is being polled.

Inactive

Indicates that polling has stopped on the component because the metric family is no longer monitored for the device.

Retired

Indicates that the component no longer exists on the physical device. Polling is stopped on the component. You can view historical data for reporting purposes. By default, retired components are not synchronized with CA Performance Center. To enable this option, select the Synchronize retired items checkbox on the Edit Data Source dialog on the Manage Data Sources page in CA Performance Center.

Filtered (interface components only)

Indicates that the component does not pass the filter criteria and polling on the component is stopped.

Note: Filtered interfaces are not displayed in CA Performance Center dashboards and data views.

5. (Optional) Select the Interface metric family and click Update Metric Family.
Data Aggregator reconfigures components for any configuration updates. For example, if you add a disk drive on a server, you can use the Update Metric Family button to rediscover the configuration update. The configuration update creates a disk component.
6. Click the Filter Report tab and follow these steps:
 - a. Look at the filters on each of the other Interface monitoring profiles to see if they are monitoring the same device collection that you want to filter.
 - b. [Remove any relationships between other Interface monitoring profiles and device collections that will block your filter criteria](#) (see page 83). For example, if your new Interface monitoring profile is associated with the All Routers device collection, remove the relationship between *other* Interface monitoring profiles and the All Routers device collection.
 - c. Run another discovery and review the updated Filter report to verify that the new filter criteria is active. If the Filter report shows that an unwanted monitoring profile was included, repeat the previous steps until you are monitoring only the interfaces that you want.

The Filter Report tab shows which interface filter criteria have been used during component monitoring. The tab also shows a report of all of the interfaces that are identified on the device and whether they matched the specified filter criteria.

Note: If you change the rules on a custom monitoring profile, the Interface Filter Criteria pane does not reflect those changes. If you disassociate the monitoring profile from a group, the Interface Filter Criteria pane does not reflect those changes. Rediscover the device to filter the interfaces that are based on the changes you made to the filter criteria and monitoring profile.

More Information:

[Discovery and Polling](#) (see page 57)

[Update Device Reconfiguration Manually](#) (see page 103)

[Clear an Interface Filter](#) (see page 116)

[Troubleshooting: Polling Has Stopped on Discovered Metric Family](#) (see page 154)

[How to Set and Activate an Interface Filter](#) (see page 115)

How to Set and Activate an Interface Filter

By default, the monitoring profile includes a filter to prevent modeling interfaces that are administratively down.

Filtering reduces the number of metric families that are monitored, which reduces unwanted data collection. For your *custom* monitoring profile, you can specify additional filters for metric families.

Note: Filters that you set on metric families are ignored when event rules that are applied to monitoring profiles trigger events.

The filter matching criteria follows the "or" rule when more than one Interface monitoring profile is assigned to a device collection. In this case, the interface that matches either filter criteria is monitored.

You can add or edit a metric family filter before or after you run a discovery. Data Aggregator applies filtering after discovery. Only component items that match the filter criteria are polled. If you add or edit a metric family filter *after* you run a discovery, polling on these metric families stops. These metric families are *not* displayed in CA Performance Center dashboards and data views.

Note: Log in as the administrator to perform this task.

To set and activate a metric family filter, follow this process:

1. If a custom monitoring profile does not exist, create a new one, or copy a profile to create a customized profile. You cannot edit or set a filter for factory monitoring profiles.
2. Select a custom monitoring profile from the Monitoring Profiles page. Click the row of a metric family from the Metric Families tab, click Edit Filter, and edit the filter criteria.

Note: Do not click directly on the metric family name because it is linked to take you to the metric family definition. Instead, click the row where the metric family name is located to activate the Edit Filter option.

- The Filter Value field is case-sensitive.
- For Speed In and Speed Out, you can use a decimal in the text field, such as 1.544 to specify an Mbps value.
- For more information about configuring Type, go to the iana web site:
<http://www.iana.org/assignments/ianaiftype-mib>
<http://www.iana.org/assignments/ianaiftype-mib>.

When you save your changes, the filter criteria appear on the Metric Families tab.

3. [Associate the monitoring profile with a device collection](#) (see page 83).
4. [Run a discovery](#) (see page 68) and then [review the Filter report on the Monitored Devices page](#) (see page 90). Look at the filters on each of the monitoring profiles to see if they are monitoring the same device collection that you want to filter.
5. [Remove any relationships between other monitoring profiles and device collections that can block your filter criteria](#) (see page 83). For example, your Interface monitoring profile could be associated with the All Routers device collection. In this case, remove the relationship between *other* monitoring profiles and the All Routers device collection.
6. Review the updated Filter report to verify that the new filter criteria are active. If the Filter report shows that an unwanted monitoring profile was included, repeat the previous steps. Eventually, no unwanted monitoring profiles are included, and you are only monitoring the metric families that you want to monitor.

Clear an Interface Filter

Interface filters can be used with custom monitoring profiles to reduce the number of interfaces that are monitored. You can clear an interface filter when you want to collect data for all device collections that are associated with a custom monitoring profile.

Note: Log in as the administrator to perform this task.

Follow these steps:

1. Navigate to the list of monitoring profiles.
2. Select a custom monitoring profile that monitors network interfaces from the list.
The Metric Families tab populates.

3. Select an Interface metric family, and click Clear Filter.

Note: This option is enabled only when you select an Interface metric family that has a filter set.

A confirmation dialog opens.

4. Click Yes.

Your changes are saved, and the filter status displays an asterisk (*) on the Metric families tab to indicate that no filters are set. The filter is applied to the next scheduled discovery (or you can manually run a discovery).

More Information:

[View Monitored Devices](#) (see page 90)

Interface Components Naming Convention

The naming convention for interface components that the Interface vendor certification or the High Speed Interface vendor certification backs is based on the following logic:

- If the `ifName` attribute exists and has a value, the interface uses this value for its name.
- If the `ifName` attribute does *not* exist or does *not* have a value, the interface uses the value of `ifDescr` for its name.

Note: New certifications for the Interface metric family can provide a different expression for the interface name.

Interface Utilization Calculation

Data Aggregator provides a means for overriding the Speed In and Speed Out values for any interface to help ensure utilization calculations use the appropriate values. For example, you could use the `bandwidth` command to configure `ifSpeedIn` and `ifSpeedOut` on your router interfaces to affect routing decisions. In this case, provide an override speed with Data Aggregator to help ensure that utilization is calculated correctly.

The settings that you make on the device can change the value to one that is higher or lower than the actual available data rate. So, the utilization calculations that are made for the interface can appear inaccurate, due to this manipulation of the bandwidth. To help ensure that interface utilization is calculated correctly, you want to provide an override speed on the interface within Data Aggregator.

Override Speed In and Speed Out Values on Interfaces

By default, utilization is calculated using the Speed In and Speed Out values that the device, which the interface is a component of, reports. However, you can override these speed values. Reporting on interface utilization can then be more accurate.

Follow these steps:

1. Click Monitored Devices from the Monitored Inventory menu for a Data Aggregator data source.

The Tree View tab displays.

2. Select Device by Collection or Device by Monitoring Profile from the drop-down list. Select the device that you want to override the Speed In and Speed Out values for an interface on and select the appropriate interface metric family on the Polled Metric Families tab.

The interface components that are monitored on the device appear in the Interface Components table.

3. Select the interface component that you want to override the Speed In and Speed Out values for and click Edit.

The Edit Interface dialog appears. The dialog displays the default discovered Speed In and Speed Out values.

4. Enter Speed In and Speed Out values in bits per second and click Save.

Note: You can remove overrides by clicking Clear and clicking Save. Going forward, bandwidth utilization charts in CA Performance Center for the interface display utilization using the speed values that the device reports. An event is generated on the interface, indicating that the speed overrides have been removed. The event can be seen in the Events Display dashboard in CA Performance Center.

The dialog closes. The overridden Speed In and Speed Out values on the interface appear in the Interface Components table with asterisks.

An event is generated on the interface, indicating that the Speed In and Speed Out values have been overridden on an interface. The event can be seen in the Events Display dashboard in CA Performance Center.

Going forward, bandwidth utilization charts in CA Performance Center for the interface display utilization using the speed values that you specified.

Chapter 6: Eventing

This section contains the following topics:

[Event Performance Guidelines](#) (see page 119)

[Performance Management Events](#) (see page 121)

[Baseline Averages](#) (see page 122)

[How to Monitor Device Performance Using Events](#) (see page 122)

[Monitoring Metrics with Event Rules](#) (see page 124)

[View Events](#) (see page 132)

[How to Configure Notifications from Event Manager](#) (see page 132)

Event Performance Guidelines

The following configuration was used to validate and benchmark event performance:

- A system in full conformance with recommended specifications for a “medium” production system of 500K polled items (referring to system sizing specifications).
- 10 event rules, spread over 7 monitoring profiles that are being used on polled items.
 - There was 1 event rule being evaluated at the 1-minute rate on a metric family comprising ~33 percent of our polled items.
 - There was 1 event rule being evaluated at the 15-minute rate on a metric family comprising ~33 percent of our polled items.
 - The remaining rules were applied to a portion of the remaining items being polled at 5 minutes.
 - The event rules were spread out evenly over 4 metric families.
 - Each rule had 1 fixed condition and 1 standard deviation condition.
 - 6 event rules had a duration of 5 minutes and window of 15 minutes.
 - 4 event rules had a duration of 15 minutes and window of 60 minutes.

Note: For optimal performance, minimize the number of monitoring profiles that have event rules for the same metric family. For example, one monitoring profile with ten rules for the Interfaces metric family will perform better than ten monitoring profiles with one rule for Interfaces metric family, when applied to the same set of devices.

- 100K polled items had a varying number of event rules that were associated to them.
- There were 5 Data Collector systems, each polling approximately 1/5th of the items.

How to Monitor Event Processing

To determine if you are doing too much eventing, you need to monitor a few key performance indicators in Data Aggregator. Eventing in Data Aggregator is performed in batches (such as, events are evaluated and generated for large groups of items at once). For this reason, we used a variety of metrics that were tracked through the Data Aggregator system's self-monitoring mechanism to assess the health of the Data Aggregator system. To view these important metrics, add a custom IM Device MultiTrend view to a dashboard. Edit the dashboard, using the following metrics from the metric family **Data Aggregator Event Calculation Times**:

- **Event Process Queue Size** – Shows the size of the event processing queue. A constant value of zero, one, or two indicates that this system is in good health and is able to maintain current eventing. A constant value larger than 2 indicates that the system is able to maintain current eventing loads, although the system is potentially behind (processing polls older than the current poll cycle). An increase in queue size without a subsequent recovery (trending downward) indicates that eventing is backed up and your system may be at risk.
- The following two metrics complement each other.
 - **Count of Cleared Events** – Number of cleared events that are in the reporting resolution window.
 - **Count of Created Events** – Number of raised events that are in the reporting resolution window.

A continuously large number of events that are raised or cleared can impact the Event Manager database.

If the combined total of these two metrics goes over 900 events in a 5-minute poll cycle, then you have exceeded the recommended 2-3 events per second generation rate recommended for medium systems. Event generation/clear bursts over the 900 events in a 5-minute poll cycle is acceptable.

- **Count of Processed Event Rule Evaluations** – An event rule evaluation is the evaluation of a single event rule against a single item. This metric tracks the sum of event rules, multiplied by the number of items those rules are applied to. The higher the number of evaluations, the more work your system is doing. However, not all evaluations are created equal. For example, evaluations with more conditions, more standard deviation conditions, or longer duration and window are more expensive than those evaluations with fewer, fixed conditions using a smaller duration and window. As such, you may be able to do more or less evaluations, depending on your event rules.

In our test environment, as described previously, we saw that exceeding 150K evaluations for a 5-minute poll cycle put the system at risk.

- **Total Time to Calculate Events** – Total amount of time that was spent processing events for this metric family. If this number exceeds the number of seconds in the reporting resolution window, then it is an indication that eventing was delayed or backlogged at that point in time.

By watching all of these metrics over time you can judge the health of event performance on your system. Additionally, if the Karaf log on the Data Aggregator system contains database and/or other errors, this can be an indication of a system under stress. In general, these self-monitored metrics should be steady. However during the evening hours (by default between 2 and 4 AM UTC), some database intensive jobs are run which can cause fluctuations in the self-monitored metrics. If the metrics return to a steady state, the system can be considered still in good health (although events can be delayed during the time the system is busy).

We recommend that you turn on eventing slowly and judge the system health before moving forward with different rules. We also recommend that you monitor the health of the system over 24 hours after each subsequent change, as there is nightly processing that can have an impact even though eventing may appear steady through-out the day-time hours.

How to Remediate When the Threshold is Exceeded

To remediate when you exceed the threshold, follow this process:

1. Turn off event rules one at a time. Check the performance after you turn off each rule before turning off another rule.
2. Reduce the number of items being polled.
3. Reduce the number of monitoring profiles with event rules that are polling items.
4. If these steps do not improve the performance, contact CA Support.

Performance Management Events

You can define two types of performance management events using event rules. You add event rules to monitoring profiles.

Time over threshold event

Is triggered by a constant (fixed-value) rule when an observed metric differs from a set fixed value for a specified duration within a window of time.

Example:

An event rule can be defined to generate an event when bandwidth utilization exceeds 80 percent for a duration of 5 minutes within a given 10-minute window of time while polling at a 5-minute interval.

Deviation from normal event

Is triggered by a standard deviation rule when an observed metric differs from what is considered to be "normal" for a specified duration within a window of time. "Normal" is based on the calculated baseline average. Initially, when limited information has been collected, the baseline average is calculated for the same hour for every day. When more data is available, Data Aggregator switches its averaging calculations to a same day of the week, same hour hourly averaging.

Example:

An event rule can be defined to generate an event when bandwidth utilization exceeds one standard deviation from the calculated same day of the week, same hour hourly average for a duration of 5 minutes within a given 10-minute window of time while polling at a 5-minute interval.

Baseline Averages

Depending on the amount of polled data that is collected, *baseline averages* are calculated in two ways:

- Initially, as an average of the hourly averages for the same hour (regardless of day).
- After enough data is collected, as an average of the hourly averages for the same day of the week, same hour.

Baseline averages help to characterize past performance for selected monitored metrics, and helps to assess present performance. Baseline averages and related standard deviations are continually calculated as each hour passes. The standard deviation provides a statistical indicator of how much variability exists in the population data that factored into the baseline average calculations.


In Data Aggregator, what is considered to be "normal" for a specified duration within a window of time is based on the calculated baseline average.

How to Monitor Device Performance Using Events

Managers (such as the Operations Center Manager) and engineers (such as an IT Operator or IT Architect), need continuous information about the health of their systems. They work with the Tool Administrator to configure Data Aggregator to generate events for devices that deviate from normal performance expectations. These events help them proactively monitor the health of their network and take remedial steps to correct a performance issue, as needed.

For example, your organization recently virtualized several business critical applications to improve efficiency. The IT Architect and the Operations Center Manager want to monitor these virtual servers to be sure they can handle the load from these applications. The Tool Administrator creates a monitoring profile and adds event rules to find over-utilized CPUs and virtual memory issues for the collection of virtual devices. Data Aggregator automatically evaluates all of the devices in the collection after every poll for each device. If needed, Data Aggregator raises or clears events when the devices satisfy event rule conditions.

The following illustration shows how to generate events automatically to help you monitor device performance issues:



As shown, the Tool Administrator works with engineers and managers to define performance expectations for a set of devices. After this discussion, the administrator decides to create a custom device collection, create a monitoring profile, and assign event rules to the monitoring profile. To begin monitoring devices, the administrator associates the monitoring profile, with its assigned event rules, to the custom device collection. As CA Performance Center generates events, the administrator, engineers, and managers can view the events in CA Performance Center.

Procedures

[Create a custom device collection](#) (see page 126).

[Add rules to a custom device collection](#) (see page 127).

[Create a monitoring profile and add event rules](#) (see page 128).

[Assign the monitoring profile to a custom device collection](#) (see page 131).

[View your events](#) (see page 132).

Monitoring Metrics with Event Rules


Events provide useful information when monitoring the health and status of your network environment. Plus, by integrating with CA Spectrum you can use events to automate processes that are based on data in an event message.

Data Aggregator handles events using monitoring profiles. Monitoring profiles can contain a set of event rules. Using metrics (from your metric families), these rules define the conditions that you want to watch for.

To implement your event rules, associate the monitoring profile with a device collection.

Important! *The key to starting and stopping the monitoring process is the device collections. Data Aggregator cannot use a monitoring profile unless you associate it to at least one device collection.*

Immediately, Data Aggregator applies the rules in that monitoring profile to the devices in that device collection. Using the metric values that are polled for these devices, the rules trigger and clear events as needed.



Events display in a CA Performance Center dashboard.

The screenshot shows the CA Performance Center interface. At the top, there's a navigation bar with 'Dashboards', 'Inventory', and 'Admin' tabs. Below this is a search bar and a 'Search All' button. The main section is titled 'Events Display'. It includes a 'Group: / All Groups [change]' dropdown and a time range selector set to 'Last 8 Hours'. Below the selector is a table of events. The table has columns: Date, Item Name, Item Type, Item Summary, Event Type, Event Status, Description, and Device Name. There are three rows of events, all dated 'Feb 25 '12 18:08 UTC'. The first two rows are about 'Cpu Load Average exceeded 2 (Maximum: 254)' and the third is about 'Percent Used exceeded 90 (Maximum: 95.396)'. At the bottom, there's a 'Search View' button and a 'Page 1 of 1' indicator.

Date	Item Name	Item T...	Item Su...	Event Type	Even...	Description	Device Name
Feb 25 '12 18:08 UTC	Sim11778:Sim61767:param04-...	Device	Server	Threshold...	Raised	Cpu Load Average exceeded 2 (Maximum: 254)	Sim11778:S...
Feb 25 '12 18:08 UTC	Sim11781:Sim11767:param04-...	Device	Server	Threshold...	Raised	Cpu Load Average exceeded 2 (Maximum: 63)	Sim11781:S...
Feb 25 '12 18:08 UTC	Sim11781:Sim11767:param04-...	Device	Server	Threshold...	Raised	Percent Used exceeded 90 (Maximum: 95.396)	Sim11781:S...

Note: You can generate user-visible alarms in CA Spectrum from events that are processed and logged in Data Aggregator. For more information, see the CA Spectrum documentation.

Create a Custom Device Collection

As the Tools Administrator for Data Aggregator, you receive a request to monitor the performance of a new group of virtual servers. The IT Architect and Operations Center Manager want to track the CPU utilization and the memory utilization. These virtual servers host critical applications, so they want frequent updates on their status.

Note: We are assuming that you already ran an initial discovery on your network and that you have discovered some virtual servers.

You decide to first create a custom device collection to group discovered virtual servers because factory (out-of-the-box) device collections do not exist for virtual servers. To create a custom device collection, you first create a custom device collection in CA Performance Center. Automatic synchronization creates a corresponding device collection in Data Aggregator.

Follow these steps:

1. Log in to CA Performance Center as a user with the Administrator role.
2. Select Admin, Custom Settings, Groups.
The Manage Groups dialog opens.
3. Right-click Monitored Collection and select Add New Group.
The Add Group dialog opens. The New tab is selected by default.
4. Supply values for the following parameters:

Group Name

Specifies a name for the group. For this example, name the group Virtual Servers.

Note: Do not use the following special characters in group names: /&\,%.

Description

(Optional) Helps you identify the group.

5. Click Save.

The Virtual Servers group appears in the Monitored groups tree. Wait for the automatic synchronization with Data Aggregator to occur. Upon synchronization, Data Aggregator creates a corresponding device collection for use in device monitoring. Synchronization can take up to 5 minutes to begin.

Add Rules to a Custom Device Collection

Networks and systems are constantly changing. Device collections are automatically updated to include devices as they are discovered. However, it can be difficult to keep custom device collections up-to-date. Therefore, you can use rules to populate the custom device collections. Newly discovered devices that meet rule specifications are added to device collections. Similarly, if they do not meet rule requirements or they are no longer monitored, devices are removed.

Group rules can be added to groups to populate and update a group contents automatically, based on various conditions. In this case, you want to add group rules to the Virtual Servers custom device collection to keep it up-to-date with discovered virtual servers. For this scenario, we assume that the IP addresses of the virtual machines are within a given range.

Follow these steps:

1. Select Admin, Custom Settings, and click Groups from the CA Performance Center main menu.

The Manage Groups dialog opens.

2. Select the group that you want to populate in the Groups tree.

Note: Devices that are added directly to a group as a manual step appear as Direct Items in the Group Properties pane. Components that are added to a group because they are children of a managed device are Inherited Items in the Group Properties.

3. Click the Rules tab, and then click Add Rule.

The Add Rule dialog opens.

4. Supply a name for the rule in the Rule Name field.

5. Select Devices from the Add list.

6. Click Add Condition.

A row of drop-down lists and fields appears.

7. Do the following actions:

- Select Device Address from the first list.
- Select 'is between' as the method for matching from the second list.
- Enter FROM *start IP address* TO *end IP address* in the third list to indicate the range where the IP addresses of the virtual machines can be found.

8. Click Preview Results to confirm that the new rule is including the devices that you want.

The results are shown in the Group Rules Preview window. You can expand each device type to see the specific devices that are added.

9. Click Save or click Save and Run Rules:
 - Save - Saves the rules without running the rules. The group is populated during the next global synchronization, which occurs approximately every 5 minutes.
 - Save and Run Rules - Saves the rules and populates the group immediately.

Create a Monitoring Profile and Add Event Rules

To set up the process of monitoring the performance of the virtual servers in your custom Virtual Servers device collection, you first create a monitoring profile and add event rules to that monitoring profile.

Event rules are not included in factory (out-of-the-box) monitoring profiles, and you cannot modify factory monitoring profiles to add event rules. You copy an existing monitoring profile to use as a basis for creating a similar profile with some changes. The changes that you will make to the custom monitoring profile will be to add event rules.

Working with the IT Architect and Operations Center Manager, you decide to create a monitoring profile and add the following event rules:

- Add a VMware memory utilization rule, as follows:
 - Violation occurs when memory utilization is above 80 percent for 300 seconds (5 minutes) within a 900-second (15-minute) window.
 - Clear the violation when the memory utilization is equal to or below 75 percent for 300 seconds within a 900-second window.
- Add a VMware CPU utilization rule, as follows:
 - Violation occurs when *both* of the following conditions are met:
 - Condition 1: CPU utilization is above 70 percent.
 - Condition 2: CPU utilization is above one standard deviation.
 - These conditions occur for 300 seconds within a 900-second window.

Follow these steps:

1. Select Admin, Data Source Settings, and click a Data Aggregator data source.
2. Click Monitoring Profiles from the Monitoring Configuration menu on the Data Aggregator administration page.

A list of monitoring profiles is populated.
3. Select the Virtual Server monitoring profile and click Copy.

The Create/Edit Monitoring Profile dialog opens.
4. Change the name of the monitoring profile to Custom Virtual Servers.

5. Click Save.

The copied monitoring profile is added to the Monitoring Profiles list.

6. Select the Custom Virtual Servers monitoring profile.
7. Click the Event Rules tab.
8. Create the VMware memory utilization event rule as follows:
 - a. Click New.

- b. Enter the following values for your new event rule:

- **Name:** VirtualMemUsageTooHigh
- **Description (optional):** VMware memory utilization
- **Metric Family:** VMware Virtual Machine
- **Duration:** 300

Note: In this example, we assume that devices are polled at the default rate of 300 seconds. The Duration value is used for the violation threshold and the clear threshold.

- **Window:** 900

Note: The Window value is used for the violation threshold and the clear threshold.

- **Severity:** Major

- c. In the 'A violation occurs when all of these conditions are met' section, select the following values:

- **Metric:** VM Memory Utilization
- **Operator:** Above
- **Value:** 80
- **Condition Type:** Fixed Value

- d. In the 'A violation is cleared when' section, select the following values:

- **Operator:** Equal to or Below
- **Value:** 75

- e. Click Save.

9. Click the Event Rules tab.
10. Create the VMware CPU utilization event rule with multiple conditions, as follows:
 - a. Click New in the Event Rules group box.

- b. Enter the following values for your new event rule:
 - **Name:** VMwareCpuUtil
 - **Description (optional):** VMware CPU utilization
 - **Metric Family:** VMware Virtual Machine
 - **Duration:** 300
 - **Window:** 900
 - **Severity:** Major
- c. In the 'A violation occurs when all of these conditions are met' section, select the following values:
 - **Metric:** CPU Utilization
 - **Operator:** Above
 - **Value:** 70
 - **Condition Type:** Fixed Value

d. Click Add Condition.

- e. In the 'A violation occurs when all of these conditions are met' section, select the following values:
 - **Metric:** CPU Utilization
 - **Operator:** Above
 - **Value:** 1
 - **Condition Type:** Standard Deviation

Note: Multiple condition event rules are limited to metrics within a single metric family. In this example, we assume that the metric family is already available for use in Data Aggregator. For more information about creating a custom metric family, see the *Data Aggregator Self-Certification Guide*.

When you define multiple conditions, the clear event is issued when any of the conditions are no longer true.

Important! It can take up to 48 hours from when you start monitoring a metric family for Data Aggregator to calculate baselines for every hour. Baseline data is needed for standard deviation rules.

11. Click Save.

Your event rules are saved. The event rules are filtered to metric families in the Custom Virtual Servers monitoring profile to help ensure all of the rules you define are evaluated.

Assign the Monitoring Profile to a Custom Device Collection

You created the Custom Virtual Servers monitoring profile and added event rules to monitor your virtual machines running critical business applications. To begin monitoring your virtual devices, and to activate event rules, assign your My Virtual Servers monitoring profile to the custom Virtual Servers device collection.

Important! *The key to starting and stopping the monitoring process is the device collections. Data Aggregator cannot use a monitoring profile unless you associate it to at least one device collection.*

Follow these steps:

1. Click Collections from the Monitoring Configuration menu on the Data Aggregator administration page.

A list of device collections displays.

2. Select the Virtual Servers device collection and click the Monitoring Profiles tab.

A list displays the monitoring profiles that are assigned to the selected devicecollection. This list will be empty.

3. Click Manage.

The Assign Collection Monitoring Profiles dialog opens.

4. Select the My Virtual Servers monitoring profile from the Available Monitoring Profiles list and click Add.

The monitoring profile moves to the Assigned Monitoring Profiles list.

5. Click Save.

Data Aggregator begins monitoring this collection of devices using your monitoring profile and event rules. Events that are generated appear in the Events Display dashboard.

View Events

CA Performance Center displays events in a report that is called the Events view. The most recent events appear first. You can control the content of the events report to display the event data most relevant to you. Features that control the report contents include the time controls and the sort and filter features.

Examples:

- **Track Configuration Changes**--When you do *not* select the Automatically Update Metric Families option on a custom monitoring profile, you must view the events log file for configuration changes and then manually click Update Metric Family in the Monitored Devices, Polled Metric Families view to help ensure that Data Aggregator picks up the device reconfiguration.
- **Troubleshoot Performance Issues**--To troubleshoot performance issues with a specific server, you can filter the events by the IP address of the server. The Events view filters the complete list of events to display only events for the selected server.

To view events, click the Dashboards menu in CA Performance Center and select "Events Display" under Operations Displays.

The Events view opens. The table displays the events that occurred within the selected time frame, listing the most recent event first.

Note: For more information about events, see the *CA Performance Center Operator Guide* and the CA Performance Center online help.

How to Configure Notifications from Event Manager

Notifications can be configured for events coming from Data Aggregator to the Event Manager. The incoming events are evaluated against the conditions that you configure for the notification criteria. Only when the criteria are met does Event Manager take a notification action. If an event does not trigger a notification, the event can still be displayed in the Event List.

The following notification types are available in the Create/Edit Notifications wizard:

Trap

Sends trap notifications to a fault or network management system (NMS) in your environment, such as CA Spectrum. Supports multiple destinations. The first destination is required.

Two MIB choices are available in the Notifications wizard to provide compatibility for customer systems.

Supported roles: Users with the Administrator role can configure trap notifications.

Email

Sends email notifications to one or more recipients when an event is raised or cleared. The email provides a link to see the context page for the device or component that triggered the alarm.

Supported roles: Users with the Create Notifications role right can configure email notifications.


A user only configures and receives notifications for events for a device in a group that the user has access to.

Consider the following information:

- Notifications are user-specific; users cannot see the notifications that other users have created.
- The Notifications option only displays when Event Manager is enabled and is in a synchronized state of Available.
- The action to delete event notifications does not affect the actual or future events.

The following diagram shows the possible workflows for the event notification options:

Event Notification Configuration Workflows



Trap or email notifications can be configured using the following process:

1. Configure event rules in the Event Rules tab of a monitoring profile in the Data Aggregator data source administration pages.
2. (Traps only) The trap receivers must be preconfigured to receive traps. Each destination can have its own configuration regarding SNMP community and IPV4 destination. For more information about trap formats, see the corresponding NMS documentation for your trap receiver.
3. (Traps only) Create an SNMP profile with the outgoing trap port (typically 162) before creating the notification.
4. (Email only) Configure the SMTP server settings by selecting Email Server from the Admin, System Settings menu in CA Performance Center.
5. Do *one* of the following actions:
 - (Administrators) Create a notification by selecting Admin, Notifications in CA Performance Center. For trap notifications, select the SNMP profile you created in step 2.

Note: As a Default Tenant Administrator, you can create a notification for a tenant administrator or tenant user by working in a real user context. Log in as a tenant administrator or tenant user. Alternatively, the Default Tenant administrator can administer the tenant and then proxy to the user to create a tenant-scoped notification.

- (Users) Create email notifications by selecting My Settings, Notifications in CA Performance Center.

Note: Administrators can also use the Event Manager API to manage notifications. Access the self-documenting interface on the Event Manager host using this URL: <http://hostname:8281/EventManager/webservice/notifications/documentation>.

Event Types

Each event created in CA Performance Management includes event type and possibly event subtype information. This information helps CA Performance Management properly process events to keep you informed about the status and health of your infrastructure.

The standard event types provided in CA Performance Management include the following information:

- Poll event—applies to events that result from polling or the analysis of poll data
- Trap event—applies to events that result from trap inputs
- Threshold event—applies to events that threshold violations on your devices trigger
- Reconfiguration change—applies to devices that are created, destroyed, or modified

- Unknown event—indicates an event of unknown type
- Any—represents a special wildcard event type that notifies subscribers about every event that is submitted to the event engine

Event types are automatically assigned upon event creation. However, CA Performance Management lets you define custom event types to help manage events consistently. Using custom event types, you can define event families that apply to your unique networking environment. When you create an event type, you determine the required attributes for that type. Thus, all events of a single event type consistently provide the same information.

After you create the custom event types, you can set up event normalization rules to map these raw events to your custom event types. With normalized events, you can ignore differences between vendors and versions when you define how to process events. Therefore, CA Performance Management can process normalized events to more specifically meet your management needs.

Note: For more information about events, see the *CA Performance Center Operator Guide*.

Chapter 7: Reporting

This section contains the following topics:

[How to Use Views](#) (see page 137)

[Baseline Averages](#) (see page 138)

[95th Percentile](#) (see page 138)

[Standard Deviation](#) (see page 139)

[Minimum and Maximum Values](#) (see page 139)

How to Use Views

Views and reports can be used in a logical workflow that lets you see the whole enterprise to identify problems, then drill down to devices and components to isolate the problem. You can also navigate directly to devices and components when you already know which system or application to troubleshoot. Information in views can also be used proactively to identify potential problems, for capacity planning, and for monthly reports about the health of the network.

Note: There can be a delay between when data is polled and when you can see the data in views and reports. If there is a data loading error, a message is logged on the Data Aggregator device.

The following types of views and reports are available:

Dashboards

Contain sets of views that let you see the polled data as meaningful information and generate reports for the top devices across the enterprise. Select Dashboards, and then select a specific dashboard from the list to open the dashboard. You can drill down to a device, and then to a device component when needed.

Device Views

Show polled data as a set of default views for a specified device. Drill down from a dashboard, or select Inventory, Devices, and then a specific device to see views of data from that device. Select a tab to see device context views in performance categories.

Device Component Views

Show multiple views simultaneously in one report for device components. Drill down from a device view, or select Inventory, Device Components, and then select a component to see the device component page.

Note: For information about customizing dashboards and views, see the CA Performance Center online help.

If data does not display in a view, drill down from the view directly to the Monitored Devices page to troubleshoot the problem. Select the **Settings** button, and click **Device Admin**. This option requires the role right named **Drill from Views into DA Admin Page**, which can be assigned to any user. The global administrator has this role right by default.

Baseline Averages

Depending on the amount of polled data that is collected, *baseline averages* are calculated in two ways:

- Initially, as an average of the hourly averages for the same hour (regardless of day).
- After enough data is collected, as an average of the hourly averages for the same day of the week, same hour.

Baseline averages help to characterize past performance for selected monitored metrics, and helps to assess present performance. Baseline averages and related standard deviations are continually calculated as each hour passes. The standard deviation provides a statistical indicator of how much variability exists in the population data that factored into the baseline average calculations.

In Data Aggregator, what is considered to be "normal" for a specified duration within a window of time is based on the calculated baseline average.

95th Percentile

A percentile is the value of a variable below which a certain percent of observations fall. For example, the 95th percentile is the value (or score) below which 95 percent of the observations are found.

95th percentile monitoring relates to bandwidth. This statistic is useful in measuring data throughput because it more accurately reflects the required capacity of the monitored link for applications that are bandwidth sensitive. The 95th percentile says that 95 percent of the time, the bandwidth usage is below this amount. The remaining 5 percent of the time, the bandwidth usage is above that amount. When using 95th percentile to perform capacity planning, we recommend setting the poll interval to at least 1-minute intervals for the monitored devices.

The 95th percentile is calculated for rollups and for reporting purposes.

A *rollup* is the process during which metric values are aggregated. In an hourly rollup, the 1 minute, 5 minute, 15 minute, 30 minute, and 60 minute polled values for metrics are aggregated every hour. In a daily rollup, hourly values for metrics are aggregated once a day. In a weekly rollup, daily values for metrics are aggregated once a week.

Standard Deviation

The *standard deviation* shows how much variation there is from the average (mean, or expected value). A low standard deviation indicates that the data points tend to be very close to the mean. High standard deviation indicates that the data points are spread out over a large range of values.

The standard deviation is calculated for rollups and for event and reporting purposes.

A *rollup* is the process during which metric values are aggregated. In an hourly rollup, the 1 minute, 5 minute, 15 minute, 30 minute, and 60 minute polled values for metrics are aggregated every hour. In a daily rollup, hourly values for metrics are aggregated once a day. In a weekly rollup, daily values for metrics are aggregated once a week.

Minimum and Maximum Values

The minimum and maximum values are calculated for rollups and for reporting purposes. These values let you observe the upper and lower bounds of performance across a given time interval.

A *rollup* is the process during which metric values are aggregated. In an hourly rollup, the 1 minute, 5 minute, 15 minute, 30 minute, and 60 minute polled values for metrics are aggregated every hour. In a daily rollup, hourly values for metrics are aggregated once a day. In a weekly rollup, daily values for metrics are aggregated once a week.

Hourly rollups:

- Minimum: The minimum value of the polled values.
- Maximum: The maximum value of the polled values.

Daily rollups:

- Minimum: The minimum value of the hourly minimums.
- Maximum: The maximum value of the hourly maximums.

Weekly rollups and beyond:

- Minimum: The minimum value of the daily minimums.
- Maximum: The maximum value of the daily maximums.

Five-minute resolution reporting:

- Minimum: The minimum value of the polled values.
- Maximum: The maximum value of the polled values.

One hour resolution reporting:

- Minimum: The minimum value of the hourly minimums.
- Maximum: The maximum value of the hourly maximums.

Day resolution reporting:

- Minimum: The minimum value of the daily minimums.
- Maximum: The maximum value of the daily maximums.

Appendix A: Calculations

This section contains the following topics:

[Baseline Average Calculations](#) (see page 141)

[95th Percentile Calculations](#) (see page 146)

[Standard Deviation Calculations](#) (see page 148)

[Total Calculations](#) (see page 150)

[Minimum and Maximum Values](#) (see page 151)

Baseline Average Calculations

Initially, when a limited amount of data has been collected, the baseline average is calculated for the same hour for every preceding day of the week. For example, after two days worth of history, a baseline average value for the 9:00 AM to 10:00 AM time period is calculated by averaging the hourly rollups for the same time periods for two consecutive days.

Eventually, when more data is available, a switchover in the calculation method occurs automatically and Data Aggregator establishes "normal" by averaging hourly samples across available preceding same days of the week. This method, then, takes into account the day of the week patterns in utilization. This method produces a better approximation of what is "normal", which can lead to a reduction in the number of missed violations and false positive events that are generated. In the same example as above, after three weeks of history, a baseline average is calculated by averaging the 9:00 AM to 10:00 AM hourly rollups for the three Mondays within the three-week period.

Note: By default, this automatic switchover occurs when at least three same day of the week, same hour data samples are available for the past 12 weeks. Data Aggregator switches back to the every day, same hour calculation method automatically when the required number of data points is no longer available. These default settings are configurable. For information about changing these default settings, see the *Data Aggregator REST Web Services Guide*.

Baseline averages are calculated for event and report generation purposes.

Example: Calculate the Same Hour Average and Population Standard Deviation for CPU Utilization

The following example shows how the "same hour" average (mean) and population standard deviation are calculated for CPU utilization on a specific device, when there are three points of data for 2:00 AM on Monday, Tuesday, and Wednesday.

Follow these steps:

1. Collect three points of data.

Day:	Monday	Tuesday	Wednesday
Mean (Average) CPU utilization:	76	65	10

2. Calculate the population mean.

The formula for calculating the population mean is as follows:

The population mean = sum of data point values in population/number of data points.

The equation for this example is as follows:

$$(76+65+10)/3$$

The population mean= 50.33

3. Calculate the difference of each data point from the mean.

The differences for this example are:

$$25.67 \quad 14.67 \quad -40.33$$

4. Calculate the square of the difference for each data point.

The squares for this example are:

$$658.78 \quad 215.11 \quad 1,626.778$$

5. Calculate the sum of the squares:

The sum of the squares for this example is 2,500.67.

6. Calculate the sum of the squares, divided by the number of data points in the population.

The result for this example is 833.56.

- The square root for this example is 28.87.

	Mean (Averages)					
	Monday	Tuesday	Wednesday	...	Mean	Std. Dev.
2:00 AM	76	65	10	...	50.33	28.87
3:00 AM	87	18	32	...	45.67	29.78
4:00 AM	10	56	40	...	35.33	19.07
5:00 AM	60	45	19	...	41.33	16.94
Hour ...	...	...	...	...	...	...

Monday of Week:	1	2	3
Mean (Averages) CPU utilization:	76	4	6

The formula for calculating the population mean is as follows:

The population mean = sum of data point values in population/number of data points.

The equation for this example is as follows:

$$(76+4+6)/3$$

The population mean = 28.67.

The differences for this example are:

47.33 -24.67 -22.67

- Calculate the square of the difference for each data point.

The squares for this example are:

2,240.44 608.44 513.78

- Calculate the sum of the squares.

The sum of the squares for this example is 3,362.67.

- Calculate the sum of the squares, divided by the number of data points in the population.

The result for this example is 1,120.89.

- Calculate the square root of the sum of squares of the data point value from the population mean.

The square root for this example is 33.48.

The standard deviation for this example is 33.48.

The following table depicts the hourly averages (mean) of rate data by day, the average (mean) of hourly averages and the population standard deviation of the hourly averages for the same day of the week, same hour:

Mean (Averages)								
	Week 1		Week 2		Week 3		Monday	
	Monday	...	Monday	...	Monday	...	Mean	Std. Dev.
2:00 AM	76	...	4	...	6	...	28.67	33.48
3:00 AM	87	...	71	...	56	...	71.33	12.66
4:00 AM	10	...	27	...	58	...	31.67	19.87
5:00 AM	60	...	3	...	32	...	31.67	23.27
Hour ...	...	...	...	...	...	...	...	...

Example: Deviation from Normal using the Same Day of the Week Same Hour Average and Population Standard Deviation for CPU Utilization

Assume that Data Aggregator is polling CPU utilization data at a 5-minute interval. You define an event rule to generate an event when CPU utilization is greater than one standard deviation above the mean for a single 5-minute poll interval.

In this example, event rule duration and window are both set to 5 minutes.

The formula for calculating when an event is raised is as follows:

$$\text{CPU utilization} = \text{mean value} + 1(\text{standard deviation value})$$

Therefore, substituting mean and standard deviation values from the preceding same day of the week, same hour for Monday at 2:00 AM is as follows:

$$\text{CPU utilization} = 28.67 + 1 (33.48)$$
$$\text{CPU utilization} = 62.15$$

As a result, if CPU utilization were to exceed 62.15 for a single 5-minute poll interval between 1:05 AM and 2:00 AM on Monday, an event would be raised. This event indicates that the CPU utilization deviated from normal for that timeframe.

Example: Examine CPU Utilization Events in a Trend Chart View


Assume that Data Aggregator is polling CPU utilization data at a 5-minute interval. In this example, you want to be alerted whenever CPU utilization on one of your business critical servers drops below the expected level. You define an event rule to generate an event when CPU utilization is one standard deviation below the mean for a single 5-minute poll interval.

For illustrative purposes only, assume that CPU utilization is 50 percent from Monday, 12:00 AM to Sunday, 12:00 AM. From Sunday, 12:00 AM to Monday, 12:00 AM, CPU utilization drops to 10 percent. You expect this drop in utilization. However, when Data Aggregator begins to calculate the baseline average, an event is raised when the CPU utilization drops to 10 percent. The event clears when the CPU utilization goes back up to 50 percent. The erroneous event is raised because, initially, when a limited amount of data is collected, the baseline average is calculated for the same hour for every day, not taking into account the difference in utilization across days of the week. Data Aggregator is expecting the CPU utilization to be 50 percent *always*.

After three weeks pass, three same days of the week, same hour data samples are available, and the baseline average calculation method changes. Data Aggregator establishes "normal" by averaging hourly samples across same days of the week. Data Aggregator is now expecting the CPU utilization to be 10 percent every Sunday at 12:00 AM to Monday at 12:00 AM. The erroneous event that was raised previously every Sunday at 12:00 AM is no longer raised.

The following view demonstrates how initially, the baseline average is calculated for the same hour for every day. When more data is available, a switchover in the calculation method occurs automatically. Data Aggregator averages hourly samples across same days of the week.

This view also demonstrates erroneous events are no longer raised when the switchover in calculation occurs.



95th Percentile Calculations

The 95th percentile is calculated for rollups, and for event and report generation purposes.

Rollups:

- For hourly rollups, the 95th percentile is calculated as a continuous percentile of the polled values.
- For daily rollups, the 95th percentile is calculated as a continuous percentile of the hourly 95th.
- For weekly rollups and beyond, the 95th percentile is calculated as a continuous percentile of the daily 95th.

Reporting:

- When the resolution is less than a day, the 95th percentile is calculated as a continuous percentile of the polled values.
- When the resolution is a day or greater, the 95th percentile is calculated as 95th of the 95th.

Example: Calculate the 95th Percentile

The following example shows how the 95th percentile is calculated, given an hour of calculation and a 5-minute poll cycle.

Follow these steps:

1. Collect an hour worth of data at a 5-minute poll cycle.

```

1  2  3  4  5  6  7  8  9 10 11 12
30 10 20 70 60 30 80 10 90 20 70 50

Reorder

10 10 20 20 30 30 50 60 70 70 80 90

```

2. Calculate the row number (RN), floor row number (FRN), and ceiling row number (CRN) values.

The formulas for calculating RN, FRN, and CRN are as follows:

- $RN = 1 + ((N - 1) * P)$

N

Represents the number of polled values that were collected.

P

Represents the percentile value.

- $FRN = \text{floor}(RN)$

FRN

Represents the largest integer that is not greater than RN.

- $CRN = \text{ceiling}(RN)$

CRN

Represents the smallest integer that is not less than RN.

The equations for this example are as follows:

$$RN = 1 + ((12 - 1) * 0.95) = 11.45$$

$$FRN = \text{floor}(RN) = 11$$

$$CRN = \text{ceiling}(RN) = 12$$

3. Calculate the 95th percentile.

The formula for calculating the 95th percentile is as follows:

if (CRN = FRN = RN) then

(value of expression from row at RN)

else

(value of expression for row at FRN) + (RN - FRN) * (CRN row - FRN row value)

The equation for this example is as follows:

$$(80) + (11.45 - 11) * (90 - 80) = 84.5000$$

The 95th percentile for this example is 84.5000.

Standard Deviation Calculations

The standard deviation is calculated for rollups and for event and report generation purposes.

Rollups:

- For hourly rollups, the standard deviation is calculated for the polled values.
- For daily rollups, the standard deviation is calculated for hourly averages.
- For weekly rollups and beyond, the standard deviation is calculated for the daily averages.

Events:

- The standard deviation provides a statistical indicator of how much variability exists in the population data that factored into the baseline average calculations.

Reporting:

- For hourly reporting, the standard deviation is calculated for the polled values.
- For daily reporting, the standard deviation is calculated for hourly averages.
- For weekly reporting and beyond, the standard deviation is calculated for the daily averages.

Example: Calculate the Standard Deviation of Population.

The following example shows how the standard deviation of the population is calculated, given 12 points of data.

The *population* refers to a set of potential values, including not only cases that are observed but those cases that are potentially observable.

The formula for calculating this standard deviation is:

population deviation = Square root of (Sum ($X - \text{population mean}$)/number of data points)

X

Is the data point value in the population.

Follow these steps:

1. Collect 12 points of data.

1	2	3	4	5	6	7	8	9	10	11	12
30	10	20	70	60	30	80	10	90	20	70	50

2. Calculate the population mean.

The population mean = sum of data point values in population/number of data points.

The population mean for this example is 45.

3. Calculate the difference of each data point from the mean.

The differences for this example are:

-15 -35 -25 25 15 -15 35 -35 45 -25 25 5

4. Calculate the square of the difference for each data point.

The squares for this example are:

225 1225 625 625 225 225 1225 1225 2025 625 625 25

5. Calculate the sum of the squares:

The sum of the squares for this example is 8900.

6. Calculate the sum of the squares, divided by the number of data points in the population.

The sum for this example is 741.6666667.

7. Calculate the square root of the sum of squares of data point value from the population mean.

The square root for this example is 27.23355773.

The standard deviation for this example is 27.23355773.

Total Calculations

The counter metric is calculated for rollups and for event and report generation purposes. The counter metric calculates the sum of all samples during a set time period. When you calculate the sum of all items in the Dynamic Trend View with a Composite Trend view type, the sum of values across all of the items that are selected in the view is calculated. On the other hand, the gauge metric type is used to calculate the average of all samples during a set time period.

Example: Calculate the Total

The following example shows how the total is calculated, given one hour of calculation and a 5-minute poll cycle.

Follow these steps:

1. Collect an hour of data at a 5-minute poll cycle.

1	2	3	4	5	6	7	8	9	10	11	12
40	10	30	60	70	20	50	20	80	30	40	60

2. Calculate the sum of the 12 samples.

The total for this example is: 510.

When considering the gauge and counter metric types, aggregation is the act of calculating the sum or average of values across all of the items or groups in a view. When you calculate the gauge across a number of aggregated items, the individual averages from the items are added. The sum of the averages is then divided by the number of items to find the gauge. Similarly, the counter is calculated by taking the individual sum of values from each item that is aggregated, and calculating the sum of all individual sums.

Example: Counter and Gauge Metrics

If you calculate the counter metric for all interfaces under a router, you can view the throughput bits. If you want to view the utilization of all interfaces, you calculate the gauge metric.

Minimum and Maximum Values

The minimum and maximum values are calculated for rollups and for reporting purposes. These values let you observe the upper and lower bounds of performance across a given time interval.

A *rollup* is the process during which metric values are aggregated. In an hourly rollup, the 1 minute, 5 minute, 15 minute, 30 minute, and 60 minute polled values for metrics are aggregated every hour. In a daily rollup, hourly values for metrics are aggregated once a day. In a weekly rollup, daily values for metrics are aggregated once a week.

Hourly rollups:

- Minimum: The minimum value of the polled values.
- Maximum: The maximum value of the polled values.

Daily rollups:

- Minimum: The minimum value of the hourly minimums.
- Maximum: The maximum value of the hourly maximums.

Weekly rollups and beyond:

- Minimum: The minimum value of the daily minimums.
- Maximum: The maximum value of the daily maximums.

Five-minute resolution reporting:

- Minimum: The minimum value of the polled values.
- Maximum: The maximum value of the polled values.

One hour resolution reporting:

- Minimum: The minimum value of the hourly minimums.
- Maximum: The maximum value of the hourly maximums.

Day resolution reporting:

- Minimum: The minimum value of the daily minimums.
- Maximum: The maximum value of the daily maximums.

Appendix B: Troubleshooting

This section contains the following topics:

[Troubleshooting: Discovery Does Not Start](#) (see page 153)

[Troubleshooting: Polling Has Stopped on Discovered Metric Family](#) (see page 154)

[Troubleshooting: Polling Stopped Event Message](#) (see page 155)

[Troubleshooting: Polling Does Not Complete for My Sensitive Device](#) (see page 155)

[Troubleshooting: Unexpected Data Aggregator Shutdown](#) (see page 156)

[Troubleshooting: I am Unable to Back Up Data Repository](#) (see page 157)

[Troubleshooting: Multiple SNMP Devices Trigger Intrusions Alarms](#) (see page 157)

Troubleshooting: Discovery Does Not Start

Symptom:

I select discovery profiles and click Run to run a discovery, but discovery fails to start, or the Run button is disabled.

Solution:

Possible reasons for a discovery failure or for a disabled Run button include the following:

- The IP domain previously specified in the discovery profile has been deleted. [Assign the discovery profile to an IP domain](#) (see page 65).
- No Data Collector has been installed for the IP domain that is specified in the selected discovery profile.

Note: For information about installing Data Collector hosts, see the *Data Aggregator Installation Guide*.

- One or more Data Collector hosts are installed for the IP domain that is specified in the selected discovery profile. However, all of the Data Collector hosts that are installed for the IP domain are stopped. [Start the Data Collector hosts](#) (see page 48).
- The tenant is deactivated. [Activate the tenant](#) (see page 99).

Troubleshooting: Polling Has Stopped on Discovered Metric Family

Symptom:

I select a device from the Monitored Devices page and see that a metric family that the device supports has stopped polling. I did not intend for polling to stop for that metric family.

Solution:

Follow this process to determine why polling has stopped and perform the appropriate steps to address the cause:

1. [Verify that a monitoring profile is defined and is set to poll the desired metric family](#) (see page 82).

If this requirement is not already met, create or edit a monitoring profile with the desired metric family defined in it.

2. [Verify that the device is associated with the device collection](#) (see page 85).

If the device is not associated with the device collection, add the device to the device collection.

Note: For information about adding a device to a device collection, see the *CA Performance Center Administrator Guide*.

3. [Verify that the monitoring profile is associated with the device collection and the device](#) (see page 82).

[If the monitoring profile is not associated, create the relationship between the monitoring profile and the device collection](#) (see page 83).

After you complete one of these actions to restart polling, select the device on the Monitored Devices page to verify:

- The status of the metric family on the Polled Metric Families tab has changed.
- The status in the Interface Components table has changed to Active.

Polling resumes automatically on existing devices.

New devices can be discovered using one of the following methods:

- [Select the polled metric family on the Monitored Devices page, and click Update Metric Family](#) (see page 90).
- Set the Change Detection rate in the monitoring profile for that metric family, with Automatic Discovery set to True.

Troubleshooting: Polling Stopped Event Message

Symptom:

A "polling stopped" event appeared in my events list. Why?

Solution:

By default, Data Aggregator controls SNMP polling, helping to ensure that too many poll requests do not overwhelm a device. One method for controlling poll traffic is the SNMP timeouts threshold. The default threshold value is 15. Therefore, when 15 or more SNMP requests timeout, polling is suspended for the remainder of the current polling cycle. An event is generated, informing you of the situation.

Note: Polling resumes at the beginning of each poll cycle. When no timeouts occur in a complete 5-minute poll cycle, a "clear" event is generated.

Troubleshooting: Polling Does Not Complete for My Sensitive Device

Symptom:

I have a critical device that I must monitor, but polling cannot complete in a single polling cycle. Sometimes there is so much network traffic that my device stops working completely. This device is known to be sensitive, but how can I reliably poll this device to help ensure good performance?

Solution:

Polling is vital for monitoring a device. However, too much polling can cause too much network traffic and can degrade your ability to monitor a device successfully. If too much network traffic is overwhelming your sensitive device, you can try the following adjustments to reduce overall traffic to the device:

- Adjust your monitoring profile to remove unnecessary metric families from polling.
- Apply a filter in your monitoring profile to reduce the number of polled interfaces.
- Adjust your monitoring profile to poll less often (for example, change the SNMP Poll Rate to 15 minutes, instead of the default 5 minutes).
- Adjust the SNMP traffic threshold to lower the number of SNMP requests that are sent to the device at a time.
- Adjust the SNMP timeouts threshold to control how many polling timeouts cause polling to suspend for the current polling cycle.

Troubleshooting: Unexpected Data Aggregator Shutdown

Symptom:

Data Aggregator shuts down unexpectedly.

Solution:

Data Aggregator shuts down if it loses contact with Data Repository. If contact with Data Repository is lost, an audit message is logged in the *Data Aggregator installation directory/apache-karaf-2.3.0/shutdown.log* file.

Note: The *Data Aggregator installation directory/apache-karaf-2.3.0/shutdown_details.log* logs heartbeat messages between Data Aggregator and Data Repository, as well as any Data Aggregator shutdowns for debugging purposes.

To resolve any connectivity or other Data Repository issues, perform the following steps:

1. Verify that the Data Repository process is running. Do the following actions:
 - a. Log in to the database server you use for Data Repository as the database administrator user, not as the root user.
 - b. Type the following command:
`/opt/vertica/bin/adminTools`
The Administration Tools dialog opens.
 - c. Select (1) View Database Cluster State.
The returning window should state: "ALL" for Host and "UP" for State.
2. If Data Repository is not running, attempt to start it by performing the following steps:
 - a. Log in to the database server you use for Data Repository.
 - b. Type the following commands:
`/opt/vertica/bin/adminTools`
The Administration Tools dialog opens.
 - c. Select (3) Start Database.
 - d. Press the Space bar next to the database name, select **OK**, and press Enter.
You are prompted for the database password.

- e. Enter the database password and press Enter.

The Data Repository database starts.

Note: If you see an error message stating that you cannot connect because of a username or password error, it is possible that a database password change is why Data Aggregator has disconnected from Data Repository.

- f. Select (E) Exit and press Enter.

If Data Repository does not start, contact CA Technical Support.

3. If Data Repository is running, you have a network connection problem, such as a network latency issue. Address your network connectivity problem.
4. Once Data Aggregator is running again, set up an automatic recovery of the Data Aggregator process.

Note: For information on setting up an automatic recovery of the Data Aggregator process, see the *Data Aggregator Installation Guide*.

Troubleshooting: I am Unable to Back Up Data Repository

Symptom:

When I run the vbr.py script to back up Data Repository, I see the message, "Another vbr instance is already running".

Solution:

This message indicates that a previous backup attempt failed, for any number of reasons (for example, password-less ssh was not set up correctly).

To reattempt to back up Data Repository, do the following steps:

1. Remove the /tmp/.initiator.mutex file from the computer where Data Repository you want to back up is installed.

The next scheduled backup will occur normally.

Troubleshooting: Multiple SNMP Devices Trigger Intrusions Alarms

Symptom:

I have many SNMP devices behind a more restricted firewall configuration (such as DMZ networks). For security reasons, the SNMP devices have different community strings. I defined an SNMP profile for each different community string, but now I am getting intrusion alarms and have been logged out of CA Performance Center.

Solution:

To find the correct SNMP profile for a device, CA Performance Center tries all of the SNMP profiles. This behavior can trigger intrusion alarms and can log you out of CA Performance Center.

To resolve this issue, follow this process:

1. Create a separate discovery profile for a critical SNMP device.
2. Assign the SNMP profile with the correct community string to the discovery profile.
3. Repeat steps one and two for each critical SNMP device.

When discovery is run, only the assigned SNMP profile is used.

Glossary

95th percentile monitoring

95th percentile monitoring relates to bandwidth. This statistic is useful in measuring data throughput because it more accurately reflects the required capacity of the monitored link for applications that are bandwidth sensitive. The 95th percentile says that 95 percent of the time, the bandwidth usage is below this amount. The remaining 5 percent of the time, the bandwidth usage is above that amount.

baseline averages

Depending on the amount of polled data that is collected, *baseline averages* are calculated in two ways:

- Initially, as an average of the hourly averages for the same hour (regardless of day).
- After enough data is collected, as an average of the hourly averages for the same day of the week, same hour.

Baseline averages help to characterize past performance for selected monitored metrics, and helps to assess present performance. Baseline averages and related standard deviations are continually calculated as each hour passes. The standard deviation provides a statistical indicator of how much variability exists in the population data that factored into the baseline average calculations.

In Data Aggregator, what is considered to be "normal" for a specified duration within a window of time is based on the calculated baseline average.

Data Collector

Data Collector coordinates data collection and actively polls for data that is used for reporting and event analysis. Operational metrics and configuration data are polled on discovered devices and their monitored components. The collected data is passed through Data Aggregator and is stored in Data Repository.

device collection

A *device collection* is a logical grouping of monitored devices, such as servers or routers.

Discovery profile

A *discovery profile* specifies how inventory discovery operates, including the IP addresses, IP address ranges, and host names that are used to locate your devices.

factory

The term "*factory*" in Data Aggregator describes items that CA Technologies provides and are often installed with the product. For example, Data Aggregator provides factory vendor certifications, monitoring profiles, and more. These out-of-the-box items can help you get Data Aggregator operational upon installation. They can also serve as examples for creating or importing custom versions of the same item. Mostly, Data Aggregator users cannot edit these factory items.

item

An *item* can be a device, component, or an interface that Data Aggregator monitors.

metric family

A *metric family* defines the set of values to collect and report on for a given technology. These values are normalized so that reporting is uniform regardless of the data source. When included in a monitoring profile, metric families determine which values to collect for the devices that are associated with that monitoring profile.

monitoring profile

A *monitoring profile* is associated with a collection of devices to specify the information to poll and the polling rate. These parameters are applied to each device in the device collection. A selection of default monitoring profiles that are based on types of devices such as routers, switches, and servers is provided.

The monitoring profile also contains the event rules that are applied to each device item in the associated device collection. Rule evaluations occur on each device item in the device collection, and on each metric that you specify in the event rules. These rule evaluations generate either raised or clear events. These events are then sent to Event Manager in CA Performance Center, CA Spectrum, and to CA Performance Center Notifier for further action.

rollup

A *rollup* is the process during which metric values are aggregated. In an hourly rollup, the 1 minute, 5 minute, 15 minute, 30 minute, and 60 minute polled values for metrics are aggregated every hour. In a daily rollup, hourly values for metrics are aggregated once a day. In a weekly rollup, daily values for metrics are aggregated once a week.

standard deviation

The *standard deviation* shows how much variation there is from the average (mean, or expected value). A low standard deviation indicates that the data points tend to be very close to the mean. High standard deviation indicates that the data points are spread out over a large range of values.