

CA Performance Management Data Aggregator

Upgrading Guide - Installation Wizard

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.

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: Upgrade Requirements and Considerations 7

Supported Upgrade Paths	7
How to Prepare for a Data Repository Upgrade	8

Chapter 2: Upgrading 9

How to Upgrade CA Performance Management Data Aggregator - Installation Wizard.....	9
Disable the Automatic Recovery of the Data Aggregator Process	10
Stop Data Collector and Data Aggregator	10
Verify the Limit on the Number of Open Files on Data Repository Hosts.....	11
Upgrade Data Repository	13
Verify the Limit on the Number of Open Files on Data Aggregator	17
Verify That All Database Tables Are Segmented	18
Segment Database Tables (Cluster Installations Only).....	19
Upgrade the Data Aggregator Installation - Installation Wizard	25
Upgrade the Data Collector Installation - Installation Wizard	30
Upgrade CA Performance Management 2.3.3 with Embedded CAMM to CA Performance Management 2.4 with CAMM 2.4.....	33
Upgrade CA Performance Management 2.3.4 with CAMM 2.2.6 to CA Performance Management 2.4.1 with CAMM 2.4	34
Re-Enable the Automatic Recovery of the Data Aggregator Process	34
Perform Post-Upgrade Steps.....	35

Chapter 3: Troubleshooting 39

Troubleshooting: Data Aggregator Synchronization Failure	39
Troubleshooting: CA Performance Center Cannot Contact Data Aggregator	40
Troubleshooting: Data Collector Installs But Does Not Appear in the Data Collector List Menu	41

Chapter 1: Upgrade Requirements and Considerations

This section contains the following topics:

[Supported Upgrade Paths](#) (see page 7)

[How to Prepare for a Data Repository Upgrade](#) (see page 8)

Supported Upgrade Paths

If you are upgrading from a previous release of Data Aggregator, upgrade your components. You always upgrade the CA Performance Center, Data Aggregator, and Data Collector components. Upgrade Data Repository when you are upgrading to the releases identified in the table that follows.

Important! If you are upgrading from Release 2.0.00 to Release 2.4, upgrade to Release 2.1.00, upgrade to Release 2.2.x, and then upgrade to Release 2.3, first.

The following table indicates the supported upgrade paths and indicates which components to upgrade:

Release	CA Performance Center Component	Data Aggregator Component	Data Collector Component	Data Repository Component
Release 2.0.00 to Release 2.1.00	Upgrade Required	Upgrade Required	Upgrade Required	Upgrade <i>Not</i> Required
Release 2.1.00 to Release 2.2.00	Upgrade Required	Upgrade Required	Upgrade Required	Upgrade Required
Release 2.2.00 to Release 2.2.1	Upgrade Required	Upgrade Required	Upgrade Required	Upgrade <i>Not</i> Required
Release 2.2.00 or 2.2.1 to 2.2.2	Upgrade Required	Upgrade Required	Upgrade Required	Upgrade Required
Release 2.2.[1, 2, 3] to 2.3.[0, 1, 2, 3]	Upgrade Required	Upgrade Required	Upgrade Required	Upgrade <i>Not</i> Required
Release 2.2.x to 2.3.4	Upgrade Required	Upgrade Required	Upgrade Required	Upgrade Required Note: Vertica Release 7 is introduced in Release 2.3.4.

Release	CA Performance Center Component	Data Aggregator Component	Data Collector Component	Data Repository Component
Release 2.3.[0, 1, 2, 3] to Release 2.3.4	Upgrade Required	Upgrade Required	Upgrade Required	Upgrade Required Note: Vertica Release 7 is introduced in Release 2.3.4.
Release 2.3.4 to Release 2.4	Upgrade Required	Upgrade Required	Upgrade Required	Upgrade <i>Not</i> Required
Release 2.3.3 to Release 2.4	Upgrade Required	Upgrade Required	Upgrade Required	Upgrade Required
Release 2.4 to Release 2.4.1	Upgrade Required	Upgrade Required	Upgrade Required	Upgrade Required
Release 2.3.4 to Release 2.4.1	Upgrade Required	Upgrade Required	Upgrade Required	Upgrade Required

Note: For information about upgrading Data Aggregator components, see the *Data Aggregator Installation Guide*. For information about upgrade requirements and considerations for releases before 2.3.x, see the *Release Notes* or *Fixed Issues* file for the release to which you are upgrading.

How to Prepare for a Data Repository Upgrade

Meet the following prerequisites before you upgrade Data Repository:

1. Be sure that you have at least 2 GB of swap space on the computer where you will install Data Repository.
2. Be sure that you are using the ext3 or ext4 file system for data and catalog directories.
3. Be sure that you are not using Logical Volume Manager (LVM) for data and catalog directories.

Chapter 2: Upgrading

This section contains the following topics:

[How to Upgrade CA Performance Management Data Aggregator - Installation Wizard](#)
(see page 9)

How to Upgrade CA Performance Management Data Aggregator - Installation Wizard

If you are upgrading from a previous release of Data Aggregator, upgrade your components.

Note: You can confirm what version of the product you have installed by looking at the .history file in the *installation_directory/logs* directory for each component.

Do the following steps in this recommended order:

1. [Disable the automatic recovery of the Data Aggregator process](#) (see page 10).
2. [Stop Data Collector and stop Data Aggregator](#) (see page 10).
3. Ensure that the user that is installing CA Performance Center has a ulimit value of at least 65536.

Note: For information about setting the ulimit value, see the *CA Performance Center Installation Guide*.

4. Upgrade CA Performance Center.

Note: For information about upgrading CA Performance Center, see the *CA Performance Center Installation Guide*.

5. [Verify the open files limit on Data Repository hosts](#) (see page 11).
6. [Upgrade Data Repository](#) (see page 13).

Note: Upgrade Data Repository only when you are upgrading to a specific release of CA Performance Management Data Aggregator. For information about supported upgrade paths and when to upgrade Data Repository, see the *Data Aggregator Release Notes*.

7. [Verify that the user that is installing Data Aggregator has a ulimit value of at least 65536](#). (see page 17)
8. [Verify that all database table projections are segmented](#) (see page 18).
9. [Segment database table projections](#) (see page 19).
10. [Upgrade Data Aggregator](#) (see page 25).

11. [Upgrade Data Collector](#) (see page 30).
12. [Re-enable the automatic recovery of the Data Aggregator process](#) (see page 34).
13. [Perform post-upgrade steps](#) (see page 35).

Disable the Automatic Recovery of the Data Aggregator Process

Disable the automatic recovery of the Data Aggregator process before you upgrade Data Aggregator. An upgrade can then be performed without having the cron job disrupt the system when it is expected to be down.

Follow these steps:

1. Log in to the computer where the Data Aggregator is installed as the root user.
2. Open a console and type the following command:

```
crontab -e
```

A vi session opens.

3. Comment out the following line:

```
* * * * * /etc/init.d/dadaemon start > /dev/null
```

For example:

```
# * * * * * /etc/init.d/dadaemon start > /dev/null
```

The automatic recovery of the Data Aggregator process is disabled.

More information:

[How to Upgrade CA Performance Management Data Aggregator - Installation Wizard](#)
(see page 9)

Stop Data Collector and Data Aggregator

Before you upgrade Data Aggregator, stop the Data Collector and Data Aggregator installations.

Follow these steps:

1. On each computer where Data Collector is installed, open a command prompt and type the following command:

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

2. On the computer where Data Aggregator is installed, open a command prompt and type the following command:

```
/etc/init.d/dadaemon stop
```

More information:

[How to Upgrade CA Performance Management Data Aggregator - Installation Wizard](#)
(see page 9)

Verify the Limit on the Number of Open Files on Data Repository Hosts

Verify that the user that is installing Data Repository has a value of at least 65536 on the number of open files. Set this value permanently.

Note: In a cluster environment, the value for the number of open files must be the same on all nodes.

Follow these steps:

1. As the root user or a sudo user, log in to each computer where you are going to install Data Repository. Open a command prompt and type the following command to verify that the number of open files is set properly:

```
ulimit -n
```

The command returns the ulimit number. This number must be at least 65536.

2. If this number is not at least 65536, do the following steps:
 - a. Open a command prompt and type the following command to change the ulimit for the open files limit to at least 65536:

```
ulimit -n ulimit_number
```

For example:

```
ulimit -n 65536
```

- b. Open the `/etc/security/limits.conf` file on each computer where you are going to install Data Repository and add the following lines:

```
# Added by Vertica
* soft nofile 65536
# Added by Vertica
* hard nofile 65536
# Added by Vertica
*      soft   fsize  unlimited
# Added by Vertica
*      hard   fsize  unlimited
```

- c. Type the following command on each computer where you are going to install Data Repository:

```
service sshd restart
```

Note: If you do not have restart as an argument, type the following commands to stop and start sshd:

```
service sshd stop
```

```
service sshd start
```

- d. To verify that the number of open files is set properly on each computer where you are going to install Data Repository, type the following command:

```
ulimit -n
```

The command returns the ulimit number that you specified earlier.

- e. (Cluster installations only) Log in as the root user or a sudo user. Ssh from one node to another node and confirm that the number of open files is set properly on each computer:

```
ulimit -n
```

The command returns the ulimit number that you specified earlier.

The user that is installing Data Repository must have a ulimit value of at least 65536. The ulimit value is set permanently. This value remains set, even if the computer where Data Repository is installed is rebooted.

- f. If the ulimit value is not set to at least 65536 on any of the hosts, do the following step on that host:

Open the `/etc/security/limits.conf` file and add the following lines:

```
# Added by Vertica
* soft nofile 65536
# Added by Vertica
* hard nofile 65536
```

The limit on the number of open files on Data Repository hosts is verified.

More information:

[How to Upgrade CA Performance Management Data Aggregator - Installation Wizard](#)
(see page 9)

Upgrade Data Repository

Upgrade Data Repository. The following scripts must be executed in sequence as part of the upgrade process:

- `dr_validate.sh` - Helps ensure that Data Repository prerequisites have been satisfied.
- `dr_install.sh` - Installs the Vertica database.

Note: For information about supported upgrade paths and when to upgrade Data Repository, see the *Data Aggregator Release Notes*.

Follow these steps:

1. Log in to the database server you use for Data Repository as the Vertica Linux database administrator user and determine what hosts Vertica is running on:
 - a. Type the following command:

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

The Administration Tools dialog opens.
 - b. Select option 6 (Configuration Menu).
 - c. Select option 3 (View Database).
 - d. Select the database.
 - e. Make a note of the hostnames. You will need these hostnames later in this procedure.
 - f. Exit the adminTools utility.
2. Open a console and log in to the computer where you plan to install Data Repository as the root user. If you require sudo installation instructions, contact CA Support.

Important! In a cluster installation, you can initiate the Data Repository installation from any of the three hosts that is participating in the cluster. The required software components are pushed to the additional two nodes during the installation.

3. Copy the `installDR.bin` file locally.
4. Change permissions for the installation file by typing the following command:

```
chmod u+x installDR.bin
```
5. To extract the installation file, type the following command:

```
./installDR.bin
```

Important! The `installDR.bin` file does not install Data Repository. This file extracts the Data Repository rpm and license file. You install Data Repository later in this procedure.

The directory that was chosen during the installDR.bin installation must be accessible by all users. **chmod** can be used to enable read/write for directories within the user home directories (for example, `chmod -R 755 ~`).

The License Agreement opens.

If you extract the Data Repository installation file from a secure shell or console and you are not running an X Window System on the computer where you want to install Data Repository, the License Agreement opens in console mode (command line). Otherwise, the agreement opens within a user interface.

6. Read the license agreement, accept the agreement, and click Next if you are in the user interface. Press Enter if you are in console mode.
7. When prompted, enter an installation directory to extract the Data Repository installation package and Vertica license file to, or accept the default installation directory of `/opt/CA/IMDataRepository_vertica7/`. Click Install and click Done if you are in the user interface. Press Enter, twice, if you are in console mode.

The Data Repository installation package and license file are extracted to the chosen directory. Three installation scripts that are required to complete the installation are extracted also.

8. To perform a manual backup of your Data Repository, type the following command:

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

Indicates the directory path and filename of the configuration file you created when you initially set up automatic backups. 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.

Important! If you back up Data Repository and have not been regularly backing up Data Repository previously, it can take several hours or more to back up Data Repository.

9. If available, copy your existing drinstall.properties from /opt/CA/IMDataRepository_vertica6 to /opt/CA/IMDataRepository_vertica7.

Note: The paths above reflect default values. Your exact locations may differ.

10. Verify that all of the parameters in the drinstall.properties file are correct. Review the following parameters:

- DbAdminLinuxUser=*The Linux user created to serve as the Vertica database administrator*

Default: dradmin

- DbAdminLinuxUserHome=*The Vertica Linux database administrator user home directory*

Default: /export/dradmin

- DbDataDir=*The location of the data directory*

Default: /data

Note: If you are unsure what your data directory is, do the following steps: Open the /opt/vertica/config/admintools.conf file. Scroll down until you see the [Nodes] section. Locate one of the lines that begins with v_dbname_nodeXXXX. This line contains the IP address of the node, the location of the catalog directory, and the location of the data directory, in that order, separated by commas. Make note of the data directory.

- DbCatalogDir=*The location of the catalog directory*

Default: /catalog

Note: If you are unsure what your catalog directory is, do the following steps: Open the /opt/vertica/config/admintools.conf file. Scroll down until you see the [Nodes] section. Locate one of the lines that begins with v_dbname_nodeXXXX. This line contains the IP address of the node, the location of the catalog directory, and the location of the data directory, in that order, separated by commas. Make note of the catalog directory.

- DbHostNames=*The comma-delimited list of hostnames for Data Repository*

Default: yourhostname1,yourhostname2,yourhostname3

- DbName=*The database name*

Default: drdata

Note: If you are unsure what your database name is, run Admintools as the Vertica Linux database administrator. From the Main Menu, select “6 Configuration Menu” and then select “3 View Database”. The name of your database is in the “Select database to view” dialog. This value should correspond with the value specified for DbName. Note the database name and select “Cancel”.

- DbPwd=*The database password*

Default: dbpass

Note: If you have an “InstallDestination” parameter in your drinstall.properties file, this parameter will no longer be used and can be safely removed.

11. Be sure that Data Repository is up and running and then type the following command to run the pre-installation script:

```
./dr_validate.sh -p properties_file
```

For example:

```
./dr_validate.sh -p drinstall.properties
```

The pre-installation script establishes passwordless SSH between all hosts in a cluster. If passwordless SSH does not exist, you are prompted for a password.

Note: The pre-installation script may prompt you to reboot.

12. Review any on-screen output for errors or warnings. You can run this script multiple times to verify that all system configuration prerequisites are set properly.
13. To run the installation script, type the following command:

```
./dr_install.sh -p properties_file
```

For example:

```
./dr_install.sh -p drinstall.properties
```

The installation script upgrades the data repository and disables unnecessary vertica processes. You may be prompted for the vertica Linux database administrator user password.

Note: Enter the password and press Enter twice to proceed.

14. Look for and resolve any failures.
15. Verify that you upgraded Data Repository correctly by doing the following steps:
 - a. Type the following command:

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

The Administration Tools dialog opens.
 - b. Verify that the top of the banner indicates that the database version is 7.0.1-2.
16. Restart Data Repository as the Vertica Linux database administrator user by selecting option 3 (Start Database) from the main menu of the Administration Tools dialog.

Data Repository is upgraded.

More information:

[How to Upgrade CA Performance Management Data Aggregator - Installation Wizard](#)
(see page 9)

Verify the Limit on the Number of Open Files on Data Aggregator

Verify that the user that is installing Data Aggregator has a value of at least 65536 on the number of open files. Set this value permanently.

Follow these steps:

1. As the root user or a sudo user, log in to the computer where you are going to install Data Aggregator. Open a command prompt and type the following command to change the ulimit for the open files limit to at least 65536:

```
ulimit -n ulimit_number
```

For example:

```
ulimit -n 65536
```

2. Open the `/etc/security/limits.conf` file on the computer where you are going to install Data Aggregator and add the following lines:

```
# Added by Data Aggregator
* soft nofile 65536
# Added by Data Aggregator
* hard nofile 65536
```

Note: Restart Data Aggregator for these changes to take effect. If you are upgrading, the upgrade process automatically restarts Data Aggregator.

3. To verify that the number of open files is set properly on the computer where you are going to install Data Aggregator, type the following command:

```
ulimit -n
```

The command returns the ulimit number that you specified earlier. The limit on the number of open files on Data Aggregator is set.

More information:

[How to Upgrade CA Performance Management Data Aggregator - Installation Wizard](#)
(see page 9)

Verify That All Database Tables Are Segmented

Verify that all database tables are segmented. Segmenting the tables reduces the amount of disk space that is required for the database. Segmenting the tables also improves general query performance.

Follow these steps:

1. As the Vertica Linux database administrator user, log in to one of the computers in the cluster where Data Repository is installed.
2. Download the `segment.py` script from where you extract the installation media. Put the script within a directory that is writeable to the Vertica Linux database administrator user. This procedure assumes that the `segment.py` script is in the home directory of the Vertica Linux database administrator user.
3. Open a command prompt and type the following command:

```
./segment.py --task tables --pass database_admin_user_password [--name database_name] [--port database_port]
```

database_admin_user_password

Indicates the Vertica Linux database administrator user password.

database_name

Indicates the name of the database. Optional, if the database name is not the default, `drdata`.

database_port

Indicates the port to use to connect to Vertica. Optional, if the port number is not the default, 5433.

For example:

```
./segment.py --task tables --pass password --name mydatabase
```

Any currently unsegmented table projections, which are sorted from largest to smallest, are returned.

4. If unsegmented database table projections are returned, [segment the tables](#) (see page 19).

More information:

[How to Upgrade CA Performance Management Data Aggregator - Installation Wizard](#)
(see page 9)

Segment Database Tables (Cluster Installations Only)

During the upgrade process, or any time after an upgrade, if you have not already done so, [verify that all database tables are segmented](#) (see page 18). If unsegmented database table projections are returned, segment them. You can also segment the database tables at any time after an upgrade.

Important! If you do not segment the database tables, you will get a warning message during the Data Aggregator component upgrade.

Segmenting the tables reduces the amount of disk space that is required for the database. Segmenting the tables also improves general query performance. You can segment the database tables either when Data Aggregator and Data Collector are up or when these components are down.

Note: Segmentation is a resource-intensive process. We *strongly* recommend that you segment the database tables when Data Aggregator and Data Collector are down, before you upgrade the Data Aggregator component. Though you can segment the database tables with Data Aggregator and Data Collector running, we advise against it.

If you segment the database tables when Data Aggregator and Data Collector are down, consider the following information before you upgrade the Data Aggregator component:

- This script can take several hours to execute large tables in the database. During an internal segmentation test and a customer database test, migration of a table 100 GB or larger took over 10 hours to complete. The segmentation time is not uniform to table size. Time depends on many factors including row count, column count, compression of the data, and machine specifications. No active monitoring of your infrastructure environment occurs when Data Aggregator and Data Collector are down.

Important! When Data Aggregator is not running, the total disk utilization during segmentation must not exceed 90 percent of available disk space. Tables which would cause the disk utilization to exceed 90 percent during segmentation will not be segmented during the process.

Consider the following information if you segment the database tables when Data Aggregator and Data Collector are running, after upgrading the Data Aggregator component:

- Do not perform any Data Aggregator administrative functions while segmenting the database tables, such as:
 - Modifying monitoring profiles
 - Associating collections to monitoring profiles
 - Increasing poll rates
 - Running new discoveries

Note: This list is not exhaustive.

- We recommend that you minimize your report load.

Important! When segmenting the tables in the database, if Data Aggregator is running, at least 40 percent of the available disk space must remain free for query processing and other database activities.

The disk space for the backup after segmentation completes will increase by the amount of data in the new segmented table projections that were created. Verify that there is enough disk space available after segmentation is completed and before backups run.

The data in the backup area for the old unsegmented table projections will be removed after the time of the restorePointLimit (entry is located in the backup configuration file) plus one day.

To avoid the time that it takes for the old data to be removed, change the snapshot name in the backup configuration file and do a full backup after segmentation is completed. You can then archive the older backup and delete the backup from the backup disk. Use the pre-segmentation backup only if you cannot use the backup that was created after segmentation completed. If you have to use the pre-segmentation backup, you will have to segment the table projections again.

Prepare for Database Table Segmentation

To prepare for database table segmentation, do the following steps:

- Backup Data Repository.
- Segment the database tables with no data.
- Estimate the amount of maintenance time that will be needed to segment the remaining database tables.

To back up Data Repository, do the following steps:

1. Back up Data Repository. Performing a backup is a time-consuming process. Execute the following command:

```
backup_script_directory_location/backup_script.sh  
>/backup_directory_location/backup.log 2>&1
```

For example:

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

Note: For information about how you originally created this script to automatically back up Data Repository, see the *CA Performance Management Administrator Guide*.

To segment the database tables with no data, do the following steps:

1. As the Vertica Linux database administrator user, log in to one of the computers in the cluster where Data Repository is installed.
2. Download the segment.py script from where you extract the installation media. Put the script within a directory that is writeable to the Vertica Linux database administrator user. This procedure assumes that the segment.py script is in the home directory of the Vertica Linux database administrator user.
3. Type the following command while Data Aggregator is running:

```
./segment.py --task zerotables --pass database_admin_user_password [--name  
database_name] [--port database_port]
```

database_admin_user_password

Indicates the Vertica Linux database administrator user password.

database_name

Indicates the name of the database. Optional, if the database name is not the default, drdata.

database_port

Indicates the port to use to connect to Vertica. Optional, if the port number is not the default, 5433.

The database tables with no data are segmented.

To determine the amount of time that you need to segment the remaining database tables, calculate a baseline:

1. To return the table names, sorted largest to smallest, type the following command:

```
./segment.py --task tables --pass database_admin_user_password [--name  
database_name] [--port database_port]
```

2. Disable scheduled backups until segmentation is complete. Backups can interfere with the segmentation process.

3. Select a table from step 1 that is about 5 GB in size. Type the following command to segment the table:

```
./segment.py --task segment --table rate_table_name --pass  
database_admin_user_password [--name database_name] [--port database_port]
```

Note: You can run this command when Data Aggregator is running, but we recommend that you run the command during a 2-3 hour maintenance window.

4. Re-enable scheduled backups.
5. Use the time it took to segment the 5 GB table to determine how long it may take to segment all of the tables that are less than 100 GB.

Note: The actual time that it takes to segment the database tables can vary based on the type and compression of the data in the tables. The values that are calculated here are rough estimates. When planning a scheduled maintenance window, add an extra hour of time for every 10 to 15 GB of database tables that will be segmented.

For large databases, you may not be able to schedule a single maintenance window that is long enough to segment the entire database. In this case, you can segment the database tables over multiple maintenance windows.

Segment Database Tables

Follow these steps:

1. As the Vertica Linux database administrator user, log in to one of the computers in the cluster where Data Repository is installed.
2. During the table projection segmentation validation in the previous procedure, if more than ten zero-length table projections were seen during this verification, type the following command to segment them:

```
./segment.py --task segment --pass database_admin_user_password --zerotables  
[--name database_name] [--port database_port]
```

database_admin_user_password

Indicates the Vertica Linux database administrator user password.

database_name

Indicates the name of the database. Optional, if the database name is not the default, drdata.

database_port

Indicates the port to use to connect to Vertica. Optional, if the port number is not the default, 5433.

For example:

```
./segment.py --task segment --pass password --zerotables --name mydatabase --port  
1122
```

3. If there are table projections that are greater than 100 GB in size, type the following command to create a script to segment the table projections that are *less than* 100 GB first:

```
./segment.py --task script --pass database_admin_user_password --lt100G [--name database_name] [--port database_port]
```

database_admin_user_password

Indicates the Vertica Linux database administrator user password.

database_name

Indicates the name of the database. Optional, if the database name is not the default, drdata.

database_port

Indicates the port to use to connect to Vertica. Optional, if the port number is not the default, 5433.

For example:

```
./segment.py --task script --pass password --lt100G --name mydatabase --port 1122
```

4. Disable scheduled backups until segmentation is complete. Backups can interfere with the segmentation process.
5. To execute the segment-script.sh script, type the following command:

```
nohup ./segment-script.sh
```

The script segments all unsegmented table projections that are less than 100 GB and sorts them from smallest to largest. The output is sent to nohup.out. If the shell is closed accidentally, the script will continue to run.

Depending on your maintenance window size and the combined size of all of the tables under 100 GB, determine which tables can be segmented in the maintenance window. Modify the generated script by removing the tables that will not fit inside the maintenance window, based on the estimated times that were calculated when you prepared for database table segmentation. Run the generated segment-script.sh during the maintenance window. If all of the tables under 100 GB could not be segmented in the maintenance window, re-generate the script and run the segment-script.sh during the next maintenance window until all of the tables have been segmented.

Important! When you run the script, any tables that will cause disk utilization to exceed 90 percent will display an error message and will not be segmented. To segment these tables, more available disk space is needed.

You will be prompted for each table that will cause disk utilization to exceed 60 percent. We strongly recommend that Data Aggregator be brought down before segmenting these tables.

Note also that this script can take several hours to execute. Do not interrupt the script execution once it begins to avoid corruption of the database.

6. Re-enable scheduled backups only if more segmentation is needed and will be done in a future maintenance window.

7. To generate a script, `segment-script.sh`, that will segment remaining table projections that are over 100 GB, type the following command:

```
./segment.py --task script --pass database_admin_user_password [--name database_name] [--port database_port]
```

database_admin_user_password

Indicates the Vertica Linux database administrator user password.

database_name

Indicates the name of the database. Optional, if the database name is not the default, `drdata`.

database_port

Indicates the port to use to connect to Vertica. Optional, if the port number is not the default, 5433.

For example:

```
./segment.py --task script --pass password --name mydatabase --port 1122
```

Important! When the script is generated, any tables that may cause disk utilization to exceed 60 percent and 90 percent are indicated.

8. Disable scheduled backups, if they are not already disabled.
9. To execute the `segment-script.sh` script, type the following command:

```
nohup ./segment-script.sh
```

The script segments all unsegmented tables and sorts them from smallest to largest.

Important! When you run the script, any tables that will cause disk utilization to exceed 90 percent will display an error message and will not be segmented. In order to segment these tables, more available disk space is needed.

You will be prompted for each table that will cause disk utilization to exceed 60 percent. We strongly recommend that Data Aggregator be brought down prior to segmenting these tables.

This script can take several hours to execute large tables in the database. During an internal segmentation test and a customer database test, segmentation of a table 100 GB or larger took over 10 hours to complete. The segmentation time is not uniform to table size. Time depends on many factors including row count, column count, compression of the data, and machine specifications. Depending on your maintenance window size, plan to segment one table per maintenance window.

10. To verify that all tables are now segmented, type the following command:

```
./segment.py --task tables --pass database_admin_user_password [--name database_name] [--port database_port]
```

The following message appears:

No tables found with unsegmented projections.

11. Re-enable scheduled backups.
12. If you segmented the database tables when Data Aggregator and Data Collector were down, start these components:
 - a. To start Data Aggregator, type the following command:

```
service dadaemon start
```
 - b. To start Data Collector, type the following command:

```
service dcmd start
```

The previous steps outline the use of the `segment.py` script, and the various things to consider when you migrate your environment. If you have any questions regarding the use of the script or if you require assistance in planning your migration, contact CA Support.

Upgrade the Data Aggregator Installation - Installation Wizard

Upgrading your existing installation of Data Aggregator lets you retain your custom profiles and configuration settings for the following features:

- Vendor certifications
- Vendor certification priorities

Note: New vendor certifications are placed at the bottom of the Vendor Certification Priorities list for the corresponding metric family. To take advantage of the new vendor certifications, manually change the vendor certification priorities. For example, F5 CPU vendor certifications are modeled as normal CPUs but do not get discovered because F5 also supports Host Resources. After an upgrade, the Host Resources CPU priority entry will be higher than the F5 entries appended to the end of the priority list. To discover F5 CPU devices and components, update the vendor certification priority for the CPU metric family. A fresh installation does not have this issue.
- Monitoring profiles
- Polling control configuration settings, as follows:
 - Compiled MIBs
 - Interface filter settings
 - Any monitored devices and components that you discovered
 - Any polled data that you gathered
 - SNMP profiles
 - Discovery profiles

The installation upgrade is installed without uninstalling the existing software. The installer detects whether there is an existing installation and confirms that you want to continue.

Important! Back up the Data Repository database before upgrading the Data Aggregator installation. Also, upgrade CA Performance Center before upgrading the Data Aggregator installation.

Follow these steps:

1. Log in to the computer where you plan to install Data Aggregator either as the root user or the sudo user.
2. Copy the installDA.bin file to the /tmp folder.
3. Change permissions for the installation file by typing the following command:

```
chmod a+x installDA.bin
```
4. To run the installation, do one of the following steps:
 - To run the installation as the root user, type the following command:

```
./installDA.bin
```
 - To run the installation as the sudo user, type the following command

```
sudo ./installDA.bin
```The License Agreement opens.
5. Read the license agreement, accept the agreement, and click Next.
6. When the installer prompts you, enter a user. This user will both own the installation and will be the user that Data Aggregator will run as.
7. Enter an installation directory when prompted.
8. The installer automatically calculates your maximum memory usage allocation for the Data Aggregator process and ActiveMQ broker. You can modify these values during or after the installation.
9. If prompted, enter the parameters for Data Repository:

Data Repository server hostname/IP

Defines either a name or an IP address for the Data Repository server host.

Note: If you installed Data Repository in a cluster, specify the IP address or name of any of the three hosts that are participating in the cluster. The installer automatically determines the name and IP address of the remaining nodes. If you are using a hardware load balancer, specify the IP address or name of the load balancer to help ensure proper connection failover handling and distribution of database activity.

Data Repository server port

Defines the port number for the Data Repository server.

Default: 5433

Database name

Defines the database name of Data Repository.

Data Repository username

Specifies the username that Data Aggregator uses to connect to the database. When installing Data Aggregator for the first time, you can specify a username and any password as long as the password does not match the username. This username and password combination is added to the database during installation.

Example: dauser

Data Repository user password

Specifies a password for the Data Repository user name.

Example: dapass

Data Repository admin username

Specify the Linux user account that was used to install Data Repository. This username is needed for administration, such as backing up and restoring Data Repository, or updating the database schema if it becomes out of synchronization. The example password that was used was dradmin.

Data Repository admin password

Defines the password for the Data Repository admin username.

Note: This database user account password was specified when you created the database after the Data Repository installation. The example password that was used was dbpassword.

10. When asked if you want the installer to recreate the schema, accept the default option. This question only applies to the case when your Data Repository has been used by a previous Data Aggregator installation.

The following table describes Data Repository users that you created:

| New User Example | Password Example | Operating System User Account? | Vertica Database User Account? |
|------------------|------------------|--------------------------------|--------------------------------|
| dauser | dapass | No | Yes |

| New User Example | Password Example | Operating System User Account? | Vertica Database User Account? |
|---|--|--------------------------------|--------------------------------|
| dradmin (This user was created during the Data Repository installation) | dbpassword
Note: The password that is specified for this database will be the password for the database administrator. | No | Yes |

The following results can occur:

- If wrong information is entered or if Data Repository is not accessible, the installer prompts the user to either correct the wrong information or choose to exit.
- If the database schema does not exist, the installer automatically creates the schema and the installation continues.
- If the database schema is out of synchronization, the installer either cancels the installation or the installer recreates the schema. The installation continues based on the options that you selected previously.
- If the database schema is correct from an earlier Data Aggregator installation, the current installation continues.

11. When prompted, enter the HTTP port number for Data Aggregator. This number is the port number for accessing Data Aggregator using the Data Aggregator REST web services and for downloading the Data Collector installer.

Default: 8581

12. When prompted, enter the SSH port for logging in to the Data Aggregator Apache Karaf shell for debugging purposes.

Default: 8501

13. Click Next.

Data Aggregator is installed.

If you chose to generate a response file when you installed Data Aggregator, a response file, `installer.properties` is created. The response file is located in the same directory where you ran the installer from. Modify the response file as needed. You can use the response file to install Data Aggregator on other computers silently.

14. To verify that the installation was successful, review the information in the file named `CA_Infrastructure_Management_Data_Aggregator_Install_timestamp.log`. This log file is located in the directory where you installed Data Aggregator, for example, `/opt/IMDataAggregator/Logs`.

15. (New installations) Register Data Aggregator as a data source with CA Performance Center.

Note: For more information about registering a data source, see the *CA Performance Center Administrator Guide*.

16. Wait a few minutes for Data Aggregator to synchronize automatically with CA Performance Center. Alternatively, you can manually synchronize CA Performance Center and Data Aggregator if you do not want to wait for the automatic synchronization to occur.

Note: The installer restarts Data Aggregator automatically when the installation is complete.

17. Be sure that Data Aggregator is up and running. Access the following address, `http://hostname:port/rest`, where *hostname:port* specifies the Data Aggregator host name and the port number. If this page displays successfully, Data Aggregator is up and running. Depending on the amount of data, Data Aggregator can take a few minutes to start running.

Note: The Karaf log on the Data Aggregator component includes the following error after upgrading the installation:

```
ERROR | tenderThread-178 | 2013-01-24 13:36:40,431 |
ndorCertificationPriorityManager | nager.core.cert-mgr.impl |
    | Failed to load the MetricFamilyVendorPriority for bundle: BundleURLEntry
[bundle=198,
resourceURL=file:/opt/IMDataAggregator/apache-karaf-2.3.0/data/cache/resource
s/198--xml-vendorpriorities-ReachabilityVendorPriorities.xml
```

18. This error for Reachability is expected and harmless. Other occurrences of this error are not expected.

More information:

[How to Upgrade CA Performance Management Data Aggregator - Installation Wizard](#)
(see page 9)

Upgrade the Data Collector Installation - Installation Wizard

You can upgrade your Data Collector installation. The installation upgrade is installed without uninstalling the existing software.

Important! Be sure that Data Aggregator is up and running before you upgrade the Data Collector installation. Access the following address, `http://hostname:port/rest`, where *hostname:port* specifies the Data Aggregator host name and the port number. If this page displays successfully, Data Aggregator is up and running.

Follow these steps:

1. Log in to the computer where you plan to install Data Collector either as the root user or the sudo user.
2. Access the Data Collector installation package by doing one of the following actions:

- If you have HTTP access to the computer where Data Aggregator is installed, open a web browser on the computer where you want to install Data Collector. Navigate to the following address and download the installation package:

`http://data_aggregator:port/dcm/install.htm`

data_aggregator:port

Specifies the Data Aggregator host name and the required port number.

Default: 8581, unless you specified a nondefault value during the Data Aggregator installation.

Save the installation package to the /tmp directory.

- If you do *not* have HTTP access to the computer where Data Aggregator is installed, open a command prompt on a computer that *does* have HTTP access. Type the following command to download the installation package to your Desktop directory:

```
wget -P /Desktop -nv  
http://data_aggregator:port/dcm/InstData/Linux/VM/install.bin
```

data_aggregator:port

Specifies the Data Aggregator host name and the required port number.

Default: 8581, unless you specified a nondefault value during the Data Aggregator installation.

Transfer the install.bin file to the /tmp directory on the computer where you want to install Data Collector.

Note: Alternatively, use the wget command when you have HTTP access to the computer where Data Aggregator is installed and want to download the Data Collector installation package in a noninteractive mode.

3. Type the following command to change to the /tmp directory:

```
cd /tmp
```

4. Change permissions for the installation file by typing the following command:

```
chmod a+x install.bin
```

5. Do one of the following actions:

- To run the installer as the root user, type the following command:

```
./install.bin
```

To run the installer as the sudo user, type the following command:

```
sudo ./install.bin
```

- To run the installer as the root user, and, at the same time, generate a response file, type the following command:

```
./install.bin -r
```

To run the installer as the sudo user, and, at the same time, generate a response file, type the following command:

```
sudo ./install.bin -r
```

6. Select your preferred language, and click OK.

The License Agreement opens.

7. Read the license agreement, accept the agreement, and click Next.

8. When the installer prompts you, enter a user. This user will both own the installation and will be the user that Data Collector will run as. The default user is the root user. Hit Enter to select the root user.

9. Enter an installation directory when prompted.

10. The installer automatically calculates your maximum memory usage allocation for the Data Collector process, basing it on 80 percent of your server memory. You can modify this value during or after the installation.

The installer prompts you for the Data Aggregator host information.

11. For a Data Aggregator to associate with Data Collector, enter either the IP address or the hostname.

Important! Specify the Data Aggregator host information correctly. If you specify the Data Aggregator host information incorrectly, Data Collector shuts down after installation. An error message is logged in the *Data Collector installation directory/apache-karaf-2.3.0/shutdown.log* file. Uninstall and reinstall Data Collector.

12. Enter either 'y' or 'n' when you are asked whether to associate this Data Collector with the Default Tenant.

Enter 'n' if you represent a service provider who is planning to deploy multi-tenancy. You can then associate each Data Collector installation with a tenant. If you are not deploying multi-tenancy, enter 'y'. For more information about multi-tenant deployments, see the CA Performance Center online help.

Data Collector is installed and started, and it connects to Data Aggregator.

Note: If you restart the computer where Data Collector is installed, Data Collector automatically restarts and connects to Data Aggregator.

If you chose to generate a response file when you installed Data Collector, a response file, `installer.properties`, is created. The response file is located in the same directory where you ran the installer. You can rename the response file. Modify the response file as needed. You can use the response file to install Data Collector on other computers silently.

13. Review the `/opt/IMDataCollector/Logs/CA_Infrastructure_Management_Data_Collector_timesamp.log` file on the computer where Data Collector is installed.

If the installation is successful, the log shows 0 Warnings, 0 NonFatalErrors, and 0 FatalErrors.

14. Verify that the Data Collector connection is successful after the installation by taking the following steps:
 - a. Log in to CA Performance Center as the global administrator or as the tenant administrator for the tenant that is associated with this Data Collector.
 - b. Navigate to the Data Aggregator administration view and expand the System Status view.
 - c. Select Data Collectors from the menu.
 - d. Verify that Data Collector appears in the list. Its Tenant and IP Domain are blank if you selected 'n' when you were asked whether to associate this Data Collector with the Default Tenant.

Note: The list can take several minutes to refresh and show the new Data Collector installation.

15. (New installations) Assign a tenant and IP domain to each Data Collector if the Tenant and IP Domain are blank:
 - a. Select the Data Collector instance and click Assign.
 - b. Select a tenant and an IP domain for this Data Collector in the Assign Data Collector dialog, and click Save.

Data Collector is installed.

When Data Collector is reinstalled, Data Collector picks up devices and components from Data Aggregator and resumes polling the devices and components that were being polled previously.

More information:

[How to Upgrade CA Performance Management Data Aggregator - Installation Wizard](#)
(see page 9)

Upgrade CA Performance Management 2.3.3 with Embedded CAMM to CA Performance Management 2.4 with CAMM 2.4

The process to upgrade from CA Performance Management 2.3.3 with embedded CA Mediation Manager (CAMM) to CA Performance Management 2.4 with CAMM 2.4 is different from upgrading CA Performance Management 2.3.4 with CAMM 2.2.6 to CA Performance Management 2.4 with CAMM 2.4. Ensure that you follow these steps to allow for a smooth transition to the most recent versions of CA Performance Management and CAMM.

Follow these steps:

1. Install CAMM 2.4 MC / LC before running the migration script. Refer to the CAMM Wiki at the following URL to learn more about installing CAMM:
<https://wiki.ca.com/display/CAMM23/Install+CA+Mediation+Manager>
<https://wiki.ca.com/display/CAMM23/Install+CA+Mediation+Manager>
Important: Make sure that you do a fresh install of CAMM 2.4.
2. After you install LC on DC, run the Device Pack Migration script
Refer to the CAMM Wiki at the following URL to learn more about migrating device packs:
<https://wiki.ca.com/display/CAMM23/Device+Pack+Migration>
<https://wiki.ca.com/display/CAMM23/Device+Pack+Migration>
Important: Make sure to run the Device Pack Migration script with the -d option. Using the -d option migrates the device packs, but does not start them automatically.
3. Stop the Data Aggregator.
4. To deploy device packs with new tags in the new directory structure for CA Performance Management 2.4, execute the `cammm-tools-cert-migration-<version>-SNAPSHOT-jar-with-dependencies.jar` script to migrate the certifications. The script is packaged with the Data Aggregator file (`CA_DA_<version>_Linux.tar.gz`). To execute the script, unzip `CA_DA_<version>_Linux.tar.gz`, which you download from the CA Performance Management page on the CA Support Site.
5. Upgrade to CA Performance Management 2.4. Refer to CA Performance Management 2.4 installation guide.
6. Start the migrated device pack engine on the CAMM MC web UI.

Upgrade CA Performance Management 2.3.4 with CAMM 2.2.6 to CA Performance Management 2.4.1 with CAMM 2.4

The process to upgrade from CA Performance Management 2.3.4 with CA Mediation Manager (CAMM) 2.2.6 to CA Performance Management 2.4.1 with CAMM 2.4 is different from upgrading CA Performance Management 2.3.3 with embedded CAMM to CA Performance Management 2.4.1 with CAMM 2.4. Ensure that you follow these steps to allow for a smooth transition to the most recent versions of CA Performance Management and CAMM.

1. Stop the Data Aggregator.
2. To deploy device packs with new tags in the new directory structure for CA Performance Management 2.4.1, execute the `camm-tools-cert-migration-<version>-SNAPSHOT-jar-with-dependencies.jar` script to migrate the certifications. The script is packaged with the Data Aggregator file (`CA_DA_<version>_Linux.tar.gz`). To execute the script, unzip `CA_DA_<version>_Linux.tar.gz`, which you download from the CA Performance Management page on the CA Support Site.
3. Upgrade to CA Performance Management 2.4.1. Refer to the CA Performance Management 2.4.1 installation guide.
4. Upgrade to CAMM 2.4. Refer to the CAMM Wiki to learn how to upgrade to CAMM 2.4 at the following URL:

<https://wiki.ca.com/display/CAMM23/Upgrade+CA+Mediation+Manager>
<https://wiki.ca.com/display/CAMM23/Upgrade+CA+Mediation+Manager>

Re-Enable the Automatic Recovery of the Data Aggregator Process

Re-enable the automatic recovery of the Data Aggregator process. You disabled the automatic recovery before you upgraded Data Aggregator. When enabled, if the database server runs out of memory, or if Data Repository is unavailable for a time, Data Aggregator shuts down automatically to help ensure that data consistency is maintained.

Follow these steps:

1. Log in to the computer where the Data Aggregator is installed as the root user.
2. Open a console and type the following command:

```
crontab -e
```

A vi session opens.

3. Uncomment out the following line by removing the pound symbol (#) from the beginning of the following line:

```
# * * * * * /etc/init.d/dadaemon start > /dev/null
```

For example:

```
* * * * * /etc/init.d/dadaemon start > /dev/null
```

The automatic recovery of the Data Aggregator process is re-enabled.

More information:

[How to Upgrade CA Performance Management Data Aggregator - Installation Wizard](#)
(see page 9)

Perform Post-Upgrade Steps

Do the following steps after you upgrade Data Aggregator:

- If you applied the Java Cryptography Extension (JCE) for Java 6, which enhances the security policy strength, consider that Data Aggregator uses Java 7. If you require this enhanced security, reapply the latest JCE after upgrading.

Note: For the Java 7 version of the JCE, visit the Oracle site.

- Upgrading Data Aggregator backs up the `/opt/IMDataAggregator/apache-karaf-X.X.X` directory to the `/opt/IMDataAggregator/backup/apache-karaf` directory. Customizations that are located in the `/opt/IMDataAggregator/apache-karaf-X.X.X/etc/` directory, such as default logging levels or other configurations are backed up, but are not restored to the installation directory automatically. To avoid losing these customizations, restore the customizations manually after you have successfully upgraded.

For example, suppose that you have a custom configuration where you update the `local-jms-broker.xml` in the `/opt/IMDataAggregator/apache-karaf-X.X.X/deploy` directory. After the upgrade, the `local-jms-broker.xml` inside the `/opt/IMDataAggregator/apache-karaf-X.X.X/deploy` directory comes from the latest installer. The customized jms broker file is backed up in the `/opt/IMDataAggregator/backup/apache-karaf` directory. To preserve the custom modification, find the backed up file and then reapply the customization to the installed directory.

- Upgrading Data Collector backs up the `/opt/IMDataCollector/apache-karaf-X.X.X` directory to the `/opt/IMDataCollector/backup/apache-karaf` directory. Customizations that are located in the `/opt/IMDataCollector/apache-karaf-X.X.X/etc/` directory, such as default logging levels or other configurations are backed up, but are not restored to the installation directory automatically. To avoid losing these customizations, restore the customizations manually after you successfully upgraded.

- Prioritize vendor certifications. After you upgrade Data Aggregator, new vendor certifications are placed at the bottom of the Vendor Certification Priorities list for the corresponding metric family. To take advantage of the new vendor certifications, manually change the vendor certification priorities. For example, F5 CPU vendor certifications are modeled as normal CPUs but do not get discovered because F5 also supports Host Resources. After an upgrade, the Host Resources CPU priority entry will be higher than the F5 entries appended to the end of the priority list. To discover F5 CPU devices and components, update the vendor certification priority for the CPU metric family.

Note: For information about prioritizing vendor certifications, see the *Data Aggregator Self-Certification Guide*.

- Reapply memory settings on CA Performance Center. For large-scale deployments, we recommend customizing your default maximum memory usage settings. These customized settings are not reapplied automatically during an upgrade. To take advantage of your custom memory settings, reapply them manually after upgrading.

Do the following steps:

1. Open *Installation*

Directory/PerformanceCenter/SERVICE/conf/wrapper.conf.old.

Note: *SERVICE* refers to the following subdirectories for the services:

- PC (Performance Center Console Service)
- DM (Performance Center Device Manager Service)
- EM (Performance Center Event Manager Service)

Example: */opt/CA/PerformanceCenter/PC/conf/wrapper.conf.old*

2. Locate the “wrapper.java.maxmemory” property, and note the specified value.
3. Open *Installation Directory/PerformanceCenter/SERVICE/conf/wrapper.conf*.
Example: */opt/CA/PerformanceCenter/PC/conf/wrapper.conf*
4. Locate the “wrapper.java.maxmemory” property, and set it to the value from step 2. Save.

Stop and restart each daemon by entering the following commands:

```
service service name stop  
service service name start
```

5. Repeat steps 1–5 for each of the services.
Your custom memory settings are reapplied.

More information:

[How to Upgrade CA Performance Management Data Aggregator - Installation Wizard](#)
(see page 9)

Chapter 3: Troubleshooting

This section contains the following topics:

[Troubleshooting: Data Aggregator Synchronization Failure](#) (see page 39)

[Troubleshooting: CA Performance Center Cannot Contact Data Aggregator](#) (see page 40)

[Troubleshooting: Data Collector Installs But Does Not Appear in the Data Collector List Menu](#) (see page 41)

Troubleshooting: Data Aggregator Synchronization Failure

Symptom:

When I try to synchronize Data Aggregator with CA Performance Center, I see a 'Synchronization failure' message. The Status column for Data Aggregator in the Manage Data Sources dialog displays 'Synchronization Failure'.

Solution:

A synchronization failure can indicate that Data Aggregator could not handle the data that was sent to it during synchronization. Review the Device Manager application log file, called DMService.log. This file appears in the CA/PerformanceCenter/DM/logs directory. The log entry shows a general SOAP exception if Data Aggregator was unable to handle data that was received from CA Performance Center during synchronization.

Look for an exception and stack trace within the following phases of synchronization:

- Pull
- Global Sync
- Bind (only executes when initially synchronizing with a data source)
- Push

Contact CA Technical Support with this information.

Troubleshooting: CA Performance Center Cannot Contact Data Aggregator

Symptom:

I installed Data Aggregator successfully, but its status in the Manage Data Sources dialog displays 'Unable to Contact.' CA Performance Center is unable to contact Data Aggregator.

Solution:

Do the following steps:

1. Log on to the Data Aggregator host computer. Open a console and type the following command to verify that Data Aggregator is running:

`service dadaemon status`
2. If Data Aggregator *is* running, a network issue is most likely preventing CA Performance Center from contacting Data Aggregator. Resolve all network problems.
3. If Data Aggregator *is not* running, start Data Aggregator. Log on to the Data Aggregator host computer as the root user or a sudo user with access to a limited set of commands. Open a console and type the following command:

`service dadaemon start`

Troubleshooting: Data Collector Installs But Does Not Appear in the Data Collector List Menu

Symptom:

I installed Data Collector successfully, but Data Collector does not appear in the Data Collector List menu.

Solution:

Do the following steps:

1. Review the *Data Collector installation directory/apache-karaf-2.3.0/shutdown.log* file to ensure that Data Collector was not shut down automatically. Data Collector is shut down automatically if you specified the Data Aggregator host, tenant, or IP domain incorrectly when you installed Data Collector. The shutdown.log file provides error information as to why Data Collector was shut down. Two main reasons why Data Collector would shut down include:
 - The Data Aggregator host information, tenant, or IP domain that was specified during the Data Collector installation were incorrect:
 - If you specified the Data Aggregator host information incorrectly, uninstall and reinstall Data Collector.
 - If you specified the tenant incorrectly, uninstall and reinstall Data Collector.
 - If you specified the IP domain incorrectly, uninstall and reinstall Data Collector.
 - Contact with Data Aggregator could not be established.
2. Type the following command to help ensure that an established connection to Data Aggregator exists:

```
netstat -a | grep 61616
```
3. If a connection to Data Aggregator does not exist, do the following steps:
 - a. View the *Data Collector installation directory/apache-karaf-2.3.0/deploy/local-jms-broker.xml* file on the Data Collector host. This file contains the hostname or IP address of the Data Aggregator host that you specified when you installed Data Collector.
 - b. Search for the “networkConnector” section of the broker.xml file. This section should contain a line as follows:

```
<networkConnector name="manager"
  uri="static:(tcp://test:61616)"
  duplex="true"
  suppressDuplicateTopicSubscriptions="false"/>
```

Ensure that the Data Aggregator hostname that is specified in the "networkConnector" section is correct and resolves through DNS or /etc/hosts entries. Data Collector cannot communicate with Data Aggregator if you entered the Data Aggregator hostname incorrectly during the Data Collector installation.

- c. Type the following command to help ensure that the connection opens successfully when you open a telnet connection to the Data Aggregator host on port 61616:

```
telnet dahostname 61616
```

This command confirms that Data Aggregator is listening in on that port.

- d. If the telnet connection does not open successfully, the reasons could be as follows:

- Data Aggregator is not running. Ensure that Data Aggregator is running. Open a console and type the following command:

```
service dadaemon status
```

If Data Aggregator is not running, start Data Aggregator. Log on to the Data Aggregator host computer as the root user or a sudo user with access to a limited set of commands. Open a console and type the following command:

```
service dadaemon start
```

- The request to initiate the connection is not making it from Data Collector to Data Aggregator successfully. Ensure that the port that is specified in the "networkConnector" section of the broker.xml file is open for incoming connections on Data Aggregator. Be sure that there are no firewall rules preventing this connection.