

CA Performance Management Data Aggregator

Installation Guide - Command Line 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: CA Performance Management Data Aggregator Installation Overview	9
Chapter 2: CA Performance Management Data Aggregator Deployment Scenarios	11
Data Aggregator Deployment Options.....	11
Small Deployment For Up To 15,000 Metrics Per Second	13
Medium Deployment For Up To 75,000 Metrics Per Second.....	15
Large Deployment For Up To 150,000 Metrics Per Second	18
Multi-tenant Deployment Considerations	20
Chapter 3: Installing	21
How to Install CA Performance Management Data Aggregator – Command Line	21
Chapter 4: Installing the Data Repository Component	23
Data Repository System Requirements.....	23
How to Prepare for a Data Repository Installation - Single Node	26
Set a Unique Hostname for the Data Repository Host.....	27
Configure the Sudo User Account for Data Repository (Optional)	28
How to Prepare for a Data Repository Installation - Cluster.....	29
Set a Unique Hostname for Each Data Repository Host	30
Configure the Sudo User Account for Data Repository (Optional)	31
Install the Data Repository Component	32
Secure Data Repository (Optional).....	37
Configure Log Rotation for Data Repository (Required)	37
How to Set Up Automatic Backups of Data Repository (Single-Node and Cluster Installations)	38
Chapter 5: Installing the Data Aggregator Component	49
Data Aggregator System Requirements	49
How to Prepare for a Data Aggregator Installation.....	52
Verify the Limit on the Number of Open Files on Data Aggregator	53
Configure the Sudo User Account for Data Aggregator (Optional)	54
Configure UTF-8 Support.....	55
Data Aggregator Considerations	56

Install Data Aggregator from the Command Line	56
---	----

Chapter 6: Installing the Data Collector Component **61**

Data Collector System Requirements	61
How to Prepare for a Data Collector Installation	64
Configure the Sudo User Account for Data Collector (Optional)	65
Configure UTF-8 Support.....	65
Set a Unique Hostname for the Data Collector Host	66
Data Collector Considerations.....	67
Install Data Collector from the Command Line	68

Chapter 7: Installing CA Performance Center **73**

Deployment Options	73
System Requirements	74
Installation Considerations.....	76
CA Performance Center Communication Ports.....	77
Linux User Account Requirements.....	78
Increase Thread Allocation in Large Deployments.....	79
Verify Time Synchronization	80
Modify Maximum Memory Usage for Each Service.....	80
Configure UTF-8 Support.....	82
Third-Party Software	83
Set the Limit on the Number of Open Files on CA Performance Center.....	83
Install CA Performance Center	84
Install Support for Non-English Languages	86

Chapter 8: Post-Installation Configuration Options **89**

How to Complete the Installation	89
Set Up Autostart on Data Repository (Optional).....	89
Configure the Automatic Recovery of the Data Aggregator Process (Recommended)	93
Modify Maximum Memory Usage for Data Aggregator and Data Collector Components After Installation (Optional)	94
Modify the External ActiveMQ Memory Limit After Installation (Optional).....	97
Change the Opened Port Number on the Data Aggregator Host (Optional)	98

Chapter 9: Troubleshooting **101**

Troubleshooting: Data Aggregator Synchronization Failure	101
Troubleshooting: CA Performance Center Cannot Contact Data Aggregator	102
Troubleshooting: Data Collector Installs But Does Not Appear in the Data Collector List Menu	103

Troubleshooting: Vertica Fails to Install in a Cluster Environment	104
--	-----

Chapter 1: CA Performance Management Data Aggregator Installation Overview

A successful Data Aggregator installation includes installing these components in the following recommended order:

1. CA Performance Center
2. Data Repository
3. Data Aggregator
4. Data Collector

Before you install any components, review the [deployment options](#) (see page 11) and decide how you want to deploy CA Performance Management in your environment. Your deployment strategy depends on the number of devices, location of these devices, and which metrics you want to monitor.

After you select a deployment plan, select an installation type:

- Single-node installation
- Cluster installation
- Sudo user installation (installing without root user access)

For each installation type, you can install using the installation wizard, command line (CLI), or silent mode. Separate guides are available for each installation method.

Chapter 2: CA Performance Management Data Aggregator Deployment Scenarios

This section contains the following topics:

[Data Aggregator Deployment Options](#) (see page 11)

[Small Deployment For Up To 15,000 Metrics Per Second](#) (see page 13)

[Medium Deployment For Up To 75,000 Metrics Per Second](#) (see page 15)

[Large Deployment For Up To 150,000 Metrics Per Second](#) (see page 18)

[Multi-tenant Deployment Considerations](#) (see page 20)


Data Aggregator Deployment Options

You have three options for a Data Aggregator deployment, depending on the amount of monitoring that you need:

- A small deployment, capable of handling 15,000 metrics per second, for monitoring up to 100,000 devices and components.
- A medium deployment, capable of handling 75,000 metrics per second, for monitoring up to 500,000 devices and components.
- A large deployment, capable of handling 150,000 metrics per second, for monitoring up to 1,000,000 devices and components.

Note: We recommend a clustered Data Repository deployment (based on three nodes) for medium and large deployments.

The following diagram demonstrates how each component works together:



Note: For Data Aggregator system requirements, see the *Data Aggregator Release Notes*.

Small Deployment For Up To 15,000 Metrics Per Second

The following information describes the requirements for a small deployment that is intended to support 15,000 metrics per second, for monitoring up to 100,000 devices and components. This deployment can be configured on the following four-host virtual machine environment:

- One CA Performance Center host
- One Data Repository host
- One Data Aggregator host
- One Data Collector host

Note: For information about CA Performance Center system requirements, see the *CA Performance Center Release Notes*.

Data Repository Requirements

The following table lists the minimum system requirements for installing the Data Repository host:

Component	Requirement
Operating system	Red Hat Enterprise Linux 5.x for x64 Red Hat Enterprise Linux 6.x for x64 Note: Upgrading an existing RH5 Vertica installation to an RH6 installation is not supported.
Processor	2 dedicated virtual CPUs
Processor speed	2.0 GHz
Memory (RAM)	16 GB
Disk space	3 partitions are required for the following directories: ■ catalog directory ■ data directory ■ backup data directory Note: For information about disk space, see the system sizing tool, which is available in the 'Recommended Reading' section on support.ca.com.

Component	Requirement
Disk input/output	200 Mega Bytes/second

Data Aggregator Requirements

The following table lists the minimum system requirements for installing the Data Aggregator host:

Component	Requirement
Operating system	Red Hat Enterprise Linux 5.x for x64 Red Hat Enterprise Linux 6.x for x64 Note: Upgrading an existing RH5 Vertica installation to an RH6 installation is not supported.
Processor	2 dedicated virtual CPUs
Processor speed	2.0 GHz
Memory (RAM)	16 GB
Disk space	50 GB Note: If Data Export will be used for one hour of data export, a second, separate 5 GB partition is required.
Disk input/output	100 Mega Bytes/second

Note: The data loading process has been rearchitected to persist all polled data temporarily to files on Data Aggregator before loading the data into Data Repository. Polled data files are removed once they are loaded into Data Repository. As a result, the amount of available disk space on Data Aggregator will fluctuate over time.

Data Collector Requirements

The following table lists the minimum system requirements for installing the Data Collector host:

Component	Requirement
Operating system	Red Hat Enterprise Linux 5.x for x64 Red Hat Enterprise Linux 6.x for x64 Note: Upgrading an existing RH5 Vertica installation to an RH6 installation is not supported.

Component	Requirement
Processor	2 dedicated virtual CPUs
Processor speed	2.0 GHz
Memory (RAM)	12 GB
Disk space	50 GB
Disk input/output	100 Mega Bytes/second

Note: For information about prerequisite steps and installation, see the *Data Aggregator Installation Guide*.

Medium Deployment For Up To 75,000 Metrics Per Second

The following information describes the requirements for a medium deployment that is intended to support 75,000 metrics per second, for monitoring up to 500,000 monitored devices and components. This deployment can be configured on the following six-host physical machine environment:

- One CA Performance Center host
- Three Data Repository hosts in a cluster
- One Data Aggregator host
- One Data Collector host

Note: For information about CA Performance Center system requirements, see the *CA Performance Center Release Notes*.

Data Repository Requirements

The following table lists the minimum system requirements for installing each Data Repository host.

Important! Although this deployment works with a single node Data Repository environment, we strongly advise a three-node cluster environment.

Component	Requirement
Operating system	Red Hat Enterprise Linux 5.x for x64 Red Hat Enterprise Linux 6.x for x64 Note: Upgrading an existing RH5 Vertica installation to an RH6 installation is not supported.
Processor	8 cores
Processor speed	2.0 GHz
Memory (RAM)	64 GB
Disk space	3 partitions are required for the following directories: ■ catalog directory ■ data directory ■ backup data directory Note: For information about disk space, see the system sizing tool, which is available in the 'Recommended Reading' section on support.ca.com.
Disk input/output	200 Mega Bytes/second

Data Aggregator Requirements

The following table lists the minimum system requirements for installing the Data Aggregator host:

Component	Requirement
Operating system	Red Hat Enterprise Linux 5.x for x64 Red Hat Enterprise Linux 6.x for x64 Note: Upgrading an existing RH5 Vertica installation to an RH6 installation is not supported.

Component	Requirement
Processor	8 cores
Processor speed	2.0 GHz
Memory (RAM)	32 GB
Disk space	50 GB Note: If Data Export will be used for one hour of data export, a second, separate 25 GB partition is required.
Disk input/output	100 Mega Bytes/second

Note: The data loading process has been rearchitected to persist all polled data temporarily to files on Data Aggregator before loading the data into Data Repository. Polled data files are removed once they are loaded into Data Repository. As a result, the amount of available disk space on Data Aggregator will fluctuate over time.

Data Collector Requirements

The following table lists the minimum system requirements for installing the Data Collector host.

Note: These requirements are for a Data Collector instance that is monitoring up to 500,000 devices and components. In practice, multiple lighter Data Collector instances can be substituted, as long as the combined capacity does not exceed 500,000 monitored devices and components. For example, you can use five virtual Data Collector instances to approximate one large one. You can see how many devices and components Data Collector is monitoring in the Data Collector List view. For more information about the Data Collector List view, see the *Data Aggregator Administrator Guide*.

Component	Requirement
Operating system	Red Hat Enterprise Linux 5.x for x64 Red Hat Enterprise Linux 6.x for x64 Note: Upgrading an existing RH5 Vertica installation to an RH6 installation is not supported.
Processor	4 cores
Processor speed	2.0 GHz
Memory (RAM)	32 GB
Disk space	50 GB
Disk input/output	100 Mega Bytes/second

Note: For information about prerequisite steps and installation, see the *Data Aggregator Installation Guide*.

Large Deployment For Up To 150,000 Metrics Per Second

The following information describes the requirements for a large deployment that is intended to support 150,000 metrics per second, for monitoring up to one million monitored devices and components. This deployment can be configured on the following seven-host physical machine environment:

- One CA Performance Center host
- Three Data Repository hosts in a cluster
- One Data Aggregator host
- Two Data Collector hosts

Note: For information about CA Performance Center system requirements, see the *CA Performance Center Release Notes*.

Data Repository Requirements

The following table lists the minimum system requirements for installing each Data Repository host.

Component	Requirement
Operating system	Red Hat Enterprise Linux 5.x for x64 Red Hat Enterprise Linux 6.x for x64 Note: Upgrading an existing RH5 Vertica installation to an RH6 installation is not supported.
Processor	12 cores
Processor speed	2.5 GHz
Memory (RAM)	96 GB

Component	Requirement
Disk space	3 partitions are required for the following directories: ■ catalog directory ■ data directory ■ backup data directory Note: For information about disk space, see the system sizing tool, which is available in the 'Recommended Reading' section on support.ca.com.
Disk input/output	200 Mega Bytes/second

Data Aggregator Requirements

The following table lists the minimum system requirements for installing the Data Aggregator host:

Component	Requirement
Operating system	Red Hat Enterprise Linux 5.x for x64 Red Hat Enterprise Linux 6.x for x64 Note: Upgrading an existing RH5 Vertica installation to an RH6 installation is not supported.
Processor	12 cores
Processor speed	2.5 GHz
Memory (RAM)	48 GB
Disk space	100 GB Note: If Data Export will be used for one hour of data export, a second, separate 5 GB partition is required.
Disk input/output	100 Mega Bytes/second

Note: The data loading process has been rearchitected to persist all polled data temporarily to files on Data Aggregator before loading the data into Data Repository. Polled data files are removed once they are loaded into Data Repository. As a result, the amount of available disk space on Data Aggregator will fluctuate over time.

Data Collector Requirements

The following table lists the minimum system requirements for installing each Data Collector host.

Important! This information is specific to *each* of the two Data Collector instances. In practice, multiple lighter Data Collector instances can be substituted, as long as the combined capacity does not exceed one million monitored devices and components. For example, you can use five virtual Data Collector instances to approximate one large one. You can see how many devices and components Data Collector is monitoring in the Data Collector List view. For more information about the Data Collector List view, see the *Data Aggregator Administrator Guide*.

Component	Requirement
Operating system	Red Hat Enterprise Linux 5.x for x64 Red Hat Enterprise Linux 6.x for x64 Note: Upgrading an existing RH5 Vertica installation to an RH6 installation is not supported.
Processor	4 cores
Processor speed	2.5 GHz
Memory (RAM)	32 GB
Disk space	50 GB
Disk input/output	100 Mega Bytes/second

Note: For information about prerequisite steps and installation, see the *Data Aggregator Installation Guide*.

Multi-tenant Deployment Considerations

In a multi-tenant deployment, note the following considerations:

- Data Aggregator *can* be shared between tenants. The information for each tenant is secure and other tenants cannot view this information.
- Data Collector is not shared among tenants. However, a tenant can have more than one Data Collector.
- Where a managed service provider is monitoring devices for multiple tenants, you can install Data Collector at the MSP site.

Note: This setup requires Data Collector to gain access through a tenant firewall to poll the devices that are being managed.

Chapter 3: Installing

This section contains the following topics:

[How to Install CA Performance Management Data Aggregator – Command Line](#) (see page 21)

How to Install CA Performance Management Data Aggregator – Command Line

You can install Data Aggregator using the command line. Installing Data Aggregator consists of installing Data Repository, Data Aggregator, and Data Collector.

Do the following steps in this recommended order:

1. Install CA Performance Center.

Note: CA Performance Center is an independent installation. It can be installed at any time. For information about installing CA Performance Center, see the *CA Performance Center Installation Guide*.

2. Install Data Repository.
3. Configure log rotation for Data Repository.
4. (Strongly Recommended) Set up automatic backups of Data Repository.
5. Install Data Aggregator.
6. (Cluster installations only) Make Data Aggregator aware of Data Repository hosts in a cluster environment.
7. Install Data Collector.

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.

Chapter 4: Installing the Data Repository Component

This section contains the following topics:

[Data Repository System Requirements](#) (see page 23)

[How to Prepare for a Data Repository Installation - Single Node](#) (see page 26)

[How to Prepare for a Data Repository Installation - Cluster](#) (see page 29)

[Install the Data Repository Component](#) (see page 32)

Data Repository System Requirements

Virtual and SAN Environments

Review the document titled *CA Support Statement for Running CA Performance Management Products in Virtualization and SAN Environments* on CA Support Online:

https://support.ca.com/phpdocs/7/5637/CA_Support_Statement_for_Running_CA_NV_M_Products_in_Virtualization_Environments.pdf

The document discusses the CA policies for installing and operating Performance Management products on virtualized servers or Storage Array Networks (SAN). You must be logged in as a registered user to view documentation on CA Support Online.

Small Deployments

A small deployment supports 15,000 metrics per second, for monitoring up to 100,000 devices and components. This deployment can be configured on the following four-host virtual machine environment:

- One CA Performance Center host
- One Data Repository host
- One Data Aggregator host
- One Data Collector host

The following information describes the minimum system requirements for installing the Data Repository host:

Component	Requirement
Operating system	Red Hat Enterprise Linux 5.x for x64 Red Hat Enterprise Linux 6.x for x64 Note: Upgrading an existing RH5 Vertica installation to an RH6 installation is not supported.
Processor	2 dedicated virtual CPUs
Processor speed	2.0 GHz
Memory (RAM)	16 GB
Disk space	3 partitions are required for the following directories: ■ catalog directory ■ data directory ■ backup data directory Note: For information about disk space, see the system sizing tool, which is available in the 'Recommended Reading' section on support.ca.com.
Disk input/output	200 Mega Bytes/second

Medium Deployments

A medium deployment supports 75,000 metrics per second, for monitoring up to 500,000 monitored devices and components. This deployment can be configured on the following six-host physical machine environment:

- One CA Performance Center host
- Three Data Repository hosts in a cluster
- One Data Aggregator host
- One Data Collector host

Important! Although this deployment works with a single-node Data Repository environment, we strongly advise a three-node cluster environment.

The following information describes the minimum system requirements for installing the Data Repository host:

Component	Requirement
Operating system	Red Hat Enterprise Linux 5.x for x64 Red Hat Enterprise Linux 6.x for x64 Note: Upgrading an existing RH5 Vertica installation to an RH6 installation is not supported.
Processor	8 cores
Processor speed	2.0 GHz
Memory (RAM)	64 GB
Disk space	3 partitions are required for the following directories: ■ catalog directory ■ data directory ■ backup data directory Note: For information about disk space, see the system sizing tool, which is available in the 'Recommended Reading' section on support.ca.com.
Disk input/output	200 Mega Bytes/second

Large Deployments

A large deployment supports 150,000 metrics per second, for monitoring up to one million monitored devices and components. This deployment can be configured on the following seven-host physical machine:

- One CA Performance Center host
- Three Data Repository hosts in a cluster
- One Data Aggregator host
- Two Data Collector hosts

The following information describes the minimum system requirements for installing the Data Repository host:

Component	Requirement
Operating system	Red Hat Enterprise Linux 5.x for x64 Red Hat Enterprise Linux 6.x for x64 Note: Upgrading an existing RH5 Vertica installation to an RH6 installation is not supported.
Processor	12 cores
Processor speed	2.5 GHz
Memory (RAM)	96 GB
Disk space	3 partitions are required for the following directories: ■ catalog directory ■ data directory ■ backup data directory Note: For information about disk space, see the system sizing tool, which is available in the 'Recommended Reading' section on support.ca.com.
Disk input/output	200 Mega Bytes/second

How to Prepare for a Data Repository Installation - Single Node

Meet the following prerequisites before you install Data Repository for a single-node cluster:

1. Verify that the following ports are open on the Data Repository systems:
 - Port 22 (TCP protocol)
 - Port 4033 (TCP and UDP protocol)
 - Port 4803 (TCP and UDP protocol)
 - Port 4804 (UDP protocol)
 - Port 4805 (UDP protocol)
 - Port 5444 (TCP protocol)
 - Port 5450 (TCP protocol)
 - Port 5433 (TCP protocol)

Note: Remote access is required to this port.
2. If a file named 'release' appears in the /etc directory, remove it. Otherwise, the Data Repository installation will fail.
3. Root access is required to install Data Repository. Determine whether you can install Data Repository as root.
4. Verify that CPU frequency scaling is disabled. Disable CPU frequency scaling through the host system Basic Input/Output System (BIOS).
5. Be sure that you have at least 2 GB of swap space on the computer where you will install Data Repository.
6. Be sure that you are using the ext3 or ext4 file system for data and catalog directories.
7. Be sure that you are not using Logical Volume Manager (LVM) for data and catalog directories.
8. [Set a unique hostname for the Data Repository host](#) (see page 30).
9. (Optional) [Configure the sudo user account](#) (see page 28).

Set a Unique Hostname for the Data Repository Host

Set a unique hostname for the computer where you plan to install Data Repository.

Follow these steps:

1. As the root user, log in to the computer where you plan to install Data Repository and verify the unique hostname on the computer.

The hostname for the computer must be associated with the IP address and *not* the loopback address of 127.0.0.1.

2. Verify that the following lines appear in the `/etc/hosts` file on the computer:

```
Do not remove the following line, or various programs
# that require network functionality will fail.
127.0.0.1 localhost.localdomain localhost
::1 localhost6.localdomain6 localhost6
IP address of your host YourHostName YourHostName.domain
```

3. If the hostname required any changes, type the following command after you make the changes:

```
service network restart
```

The `/etc/hosts` file is configured correctly.

The unique host name is set.

More information:

[How to Prepare for a Data Repository Installation - Single Node](#) (see page 26)

[How to Prepare for a Data Repository Installation - Cluster](#) (see page 29)

Configure the Sudo User Account for Data Repository (Optional)

Before you install Data Repository, log in as the root user. However, in some environments, unrestricted root user access is not available. If root user access is not available, a sudo user with access to a limited set of commands can install and run the software.

Follow these steps:

1. Log in to the computer where you want to install Data Repository as the root user.
2. Add the following command alias to the command alias section of the `/etc/sudoers` file:

```
Cmnd_Alias CA_DATAREP =
/tmp/installDR.bin,/opt/CA/IMDataRepository_vertica7/dr_validate.sh,/opt/CA/IMDataRepository_vertica7/dr_install.sh,/usr/bin/vim,/usr/bin/reboot
```

```
## Allows the Data Repository user to manage the Data Repository
```

```
dasudouser_name ALL = CA_DATAREP
```

This command alias details the commands that the sudo user must be able to run.

The sudo user account is configured.

More information:

[How to Prepare for a Data Repository Installation - Single Node](#) (see page 26)

[How to Prepare for a Data Repository Installation - Cluster](#) (see page 29)

How to Prepare for a Data Repository Installation - Cluster

Meet the following prerequisites before you install Data Repository for a cluster:

1. Verify that the following ports are open on the Data Repository systems:

- Port 22 (TCP protocol)
- Port 4033 (TCP and UDP protocol)
- Port 4803 (TCP and UDP protocol)
- Port 4804 (UDP protocol)
- Port 4805 (UDP protocol)
- Port 5444 (TCP protocol)
- Port 5450 (TCP protocol)
- Port 5433 (TCP protocol)

Note: Remote access is required to this port.

- Ports 48073 and higher must be open for intercluster communication.

2. When installing Data Repository in a cluster, select the hosts where you will install Data Repository nodes.

Important! Database software is deployed on each participating host in a cluster. This software represents a 'node' in the cluster. A three-node cluster represents the simplest configuration that can tolerate the loss of a single node. You can, however, include more than three hosts in the cluster. If more than one node fails or shuts down, Data Repository is no longer available for use and Data Aggregator shuts down automatically.

3. If a file named 'release' appears in the /etc directory, remove it. Otherwise, the Data Repository installation will fail.
4. Verify that the root user or sudo user can create database administrator user accounts, or have an administrator create these accounts.
5. Verify that CPU frequency scaling is disabled. Disable CPU frequency scaling through the host system Basic Input/Output System (BIOS).
6. Verify all of the hosts in the cluster are in the same subnet.

7. Verify that the root user can use Secure Shell (SSH) to log in (ssh) to all of the hosts in the cluster.
8. Be sure that you have at least 2 GB of swap space on each computer where you will install Data Repository.
9. Be sure that you are using the ext3 or ext4 file system for data and catalog directories.
10. Be sure that you are not using Logical Volume Manager (LVM) for data and catalog directories.
11. [Set a unique hostname for each Data Repository host](#) (see page 30).
12. (Optional) [Configure the sudo user account](#) (see page 28).

Set a Unique Hostname for Each Data Repository Host

Set a unique hostname for each computer where you plan to install Data Repository. Three unique hostnames are required in a cluster installation.

Follow these steps:

1. As the root user, log in to each computer where you plan to install Data Repository and verify the unique hostname on each computer.

The hostname for each computer must be associated with the IP address and *not* the loopback address of 127.0.0.1.

2. Verify that the following lines appear in the /etc/hosts file on each computer:

```
Do not remove the following line, or various programs
# that require network functionality will fail.
127.0.0.1 localhost.localdomain localhost
::1 localhost6.localdomain6 localhost6
IP address of your host YourHostName YourHostName.domain
```

3. If the hostname required any changes, type the following command after you make the changes:

```
service network restart
```

The /etc/hosts file is configured correctly.

The unique host name is set.

4. (Cluster installations only) The hostnames of all hosts in the cluster must resolve correctly. If the hostname resolution is incorrect, the Data Repository cluster does not install or work properly. All participating hosts in the cluster must use static IP or permanently-leased DHCP addresses. Set up the /etc/hosts file on each of the three hosts you selected for the cluster. The hosts file must contain entries for all three hosts in the cluster.

For example, if the hosts in the cluster are named host01, host02, and host03, the `/etc/hosts` file on each host resembles the following example:

```
127.0.0.1 localhost.localdomain localhost
192.168.13.128 host01.domain host01
192.168.13.129 host02.domain host02
192.168.13.130 host03.domain host03
```

Note: Do not remove the loopback address (127.0.0.1) line. The local Data Repository hostname cannot be on the 127.0.0.1 line. Also, do not use the loopback address or localhost name when you are defining hosts in the cluster.

5. Verify that hostname resolution works for each host in the cluster.

For example, on host01, the following syntax is correct:

```
$ /bin/hostname -f
host01
```

Hostname resolution is configured.

More information:

[How to Prepare for a Data Repository Installation - Single Node](#) (see page 26)

[How to Prepare for a Data Repository Installation - Cluster](#) (see page 29)

Configure the Sudo User Account for Data Repository (Optional)

Before you install Data Repository, log in as the root user. However, in some environments, unrestricted root user access is not available. If root user access is not available, a sudo user with access to a limited set of commands can install and run the software.

Follow these steps:

1. Log in to the computer where you want to install Data Repository as the root user.
2. Add the following command alias to the command alias section of the `/etc/sudoers` file:

```
Cmnd_Alias CA_DATAREP =
/tmp/installDR.bin,/opt/CA/IMDataRepository_vertica7/dr_validate.sh,/opt/CA/I
MDataRepository_vertica7/dr_install.sh,/usr/bin/vim,/usr/bin/reboot
```

```
## Allows the Data Repository user to manage the Data Repository
```

```
dasudouser_name ALL = CA_DATAREP
```

This command alias details the commands that the sudo user must be able to run.

The sudo user account is configured.

More information:

[How to Prepare for a Data Repository Installation - Single Node](#) (see page 26)

[How to Prepare for a Data Repository Installation - Cluster](#) (see page 29)

Install the Data Repository Component

After you meet the prerequisites, you can install Data Repository. Install Data Repository before you install Data Aggregator. The following scripts must be executed in sequence as part of the installation process:

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

Each script, when run, generates a corresponding log file in the *installation_directory/logs* directory on the Data Repository host from which the scripts were run. These log files include the step-by-step output of the scripts. You can review the script outputs to validate successful/failed script runs.

Important! Before you install Data Repository, review the system requirements.

Follow these steps:

1. Open a console and log in to the computer where you plan to install Data Repository as the root user.

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.

2. Copy the `installDR.bin` file locally. For the sake of installation instructions, we assume that the installer is placed in the `/tmp` directory, however, this location is not required.
3. Change permissions for the installation file by typing the following command:

```
chmod u+x installDR.bin
```

4. To extract the installation file, do one of the following steps:
 - To extract the installation file as the root user, type the following command:

```
./installDR.bin
```

- To extract the installation file as the sudo user, type the following command:

```
sudo ./installDR.
```

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

The License Agreement opens.

If you extract the Data Repository installation file from a secure shell or console without running an X Window System on the computer on which you install Data Repository, the License Agreement opens in console mode (command line). Otherwise, the agreement opens within a user interface.

5. 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.
6. 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. If you are in console mode, press Enter twice.

Note: Do not use the Logical Volume Manager (LVM) for the /opt directory.

The Data Repository installation package, license file, and associated setup scripts are extracted to the chosen directory.

7. Adjust the following parameters in the drinstall.properties file to reflect your installation-specific values. The drinstall.properties file exists in the installation directory you specified previously.

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

Default: dradmin

Note: If this user is not found in the system, the Vertica installer creates it. This user is the Vertica database administrator. If the dradmin user is not created, changing to user Vertica automatically creates the dradmin user.

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

Default: /export/dradmin

Note: This directory is created if the Vertica installer creates the user. Be sure that the directory leading up to the home account already exists on the system. For example, if you are using /export/dradmin, be sure that /export exists.

- DbDataDir=*The location of the data directory*

Default: /data

Note: Do not use the Logical Volume Manager (LVM) for the data directory.

- DbCatalogDir=*The location of the catalog directory*

Default: /catalog

Note: Do not use the Logical Volume Manager (LVM) for the catalog directory.

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

Default: yourhostname1,yourhostname2,yourhostname3

- DbName=*The database name*

Default: drdata

- DbPwd=*The database password*

Default: dbpass

Note: The database password that you define here is used during the installation of Data Aggregator.

8. (Optional) Set up passwordless SSH for the root user in cluster environments from one Data Repository host to another:

- a. Open a console and log in to the Data Repository host as the root 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 root user public key into the remote host's list of authorized keys, type the following command:

```
ssh-copy-id -i root_user@remotehost
```

remotehost

Is another host in the cluster where you are copying the SSH ID.

- d. To verify that passwordless ssh is set up correctly, log in to the remote host from the local host:

```
ssh root_user@remotehost ls
```

- e. Repeat steps 1-4 for each pair of hosts.

Note: A three-node cluster requires six variations of the previous steps.

If the passwordless SSH has been set up successfully, you are not prompted for a password. You also see a directory listing from the 'ls command'.

9. To run the validation script, do one of the following steps:

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

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

For example:

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

- To run the validation script as the sudo user, type the following command:

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

For example:

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

Note: If you run the validation script as the sudo user, you are prompted for the Vertica database administrator password. Sometimes, you are prompted multiple times.

The validation script establishes SSH without a password for the root user across all hosts in a cluster. If SSH without a password does not exist for the root account, you are prompted for a password. You are sometimes prompted multiple times.

Note: The validation script sometimes asks you to reboot.

10. Review any on-screen output for failures or warnings. You can run this script multiple times after you fix any failures or warnings. The script automatically corrects many failures or warnings. Proceed only if the final status is "PASSED". If the final status is not "PASSED", contact CA Support.

11. To run the installation script, do one of the following steps:

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

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

- To run the installation script as the sudo user, type the following command:

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

The installation script installs the data repository, creates the database, and disables unnecessary Vertica processes. If the database administrator user does not already exist, the installation script also creates it. The script then prompts you to assign a new password.

12. Look for and resolve any failures.

13. Verify that Data Repository has been installed successfully by doing the following steps:

- a. To log in to the database server you use for Data Repository as the database administrator user, type the following command:

```
su - dradmin
```

- b. Type the following command:

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

- c. The Administration Tools dialog opens.

- d. Select (1) View Database Cluster State and then select OK or press Enter.

The database name appears and the State is reported as UP.

- e. Select OK to acknowledge that the database is UP.
- f. Select (E) Exit and press Enter.

Note: If the database does not start automatically, select Start DB to start the database manually. If the database is not started, the Data Aggregator installation fails.

- 14. (Optional) [Secure Data Repository](#) (see page 37).
- 15. (Required) [Configure the log rotation for Data Repository](#) (see page 37).
- 16. (Strongly Recommended) [Set up automatic backups](#) (see page 38).

When Data Repository is installed, three users are created:

New User Example	Password Example	Operating System User Account?	Vertica Database User Account?	Notes	Permissions
spread	N/A	Yes	No	This user is an internal user that Vertica creates. Do not do anything with this user.	This daemon-only user owns the Data Repository processes.
dradmin	drpass	Yes	No	This user is the first user that you created when you installed Data Repository. When the dradmin user is created, a verticadba group is also created. The dradmin user is added to this group.	This user can run the Data Repository processes and the Administration Tools utility. This user owns Data Repository catalog files, data files, and so on.
dradmin Note: This user is different from the user that is displayed in the previous row.	dbpassword Note: The password that is specified for this database is the password for the database administrator.	No	Yes	This user was created when the database was created.	

Note: Vertica includes a `verticadba` group for tighter control over filesystem access in the `/opt/vertica/` directories. During the installation, the `verticadba` group is created, and existing users are added to the group with permissions set to 775. This setting grants full privileges to the `verticadba` group and read/execute privileges to all other users. The `/opt/vertica/log` and `/opt/vertica/config` directories are the folders with the modified permissions.

Secure Data Repository (Optional)

If you want to limit the users who can log in to the database to just the Data Repository administrative account and the root user, lock down the database.

Follow these steps:

1. Modify the `/etc/pam.d/sshd` file by adding the following entry, for the PAM access module, after the "account required pam_nologin.so" entry:

```
account required pam_access.so accessfile=/etc/security/sshd.conf
```

2. Remove the following line from the `/etc/security/access.conf` file:

```
-:ALL EXCEPT database_admin_user root:LOCAL
```

For example:

```
-:ALL EXCEPT dradmin root:LOCAL
```

More information:

[Install the Data Repository Component](#) (see page 32)

Configure Log Rotation for Data Repository (Required)

To prevent the underlying `vertica.log` file from becoming too large, configure log rotation for Data Repository. The recommended configuration for the log rotation is a daily rotation with logs retained for 21 days.

Important! Configuring the log rotation is required because the underlying Data Repository log file (`vertica.log`) can grow substantially.

Follow these steps:

1. Log in to the database server for Data Repository as the database administrator user. Type the following command:

```
su - dradmin
```

2. Type the following command:

```
/opt/vertica/bin/admintools -t logrotate -d drdata -r daily -k 3
```

- -d indicates the database name.
- -r indicates how often to rotate the logs (daily, weekly, monthly).
- -k indicates the number of weeks to keep the log.

Data Repository log rotation is now set up to occur daily, and only three weeks of log files are maintained.

You can verify that the vertica.log rotation has been configured correctly. As the log rotation occurs, new gzipped vertica.log files appear in the Vertica catalog directory for previous days. These log files have filenames such as vertica.log.1.gz, vertica.log.2.gz, and so on, with vertica.log.1.gz being the most recent backup.

More Information:

[Install the Data Repository Component](#) (see page 32)

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 39).
2. Do one of the following steps:
 - [Configure a Data Repository backup to a remote host](#) (see page 40).
 - [Configure a Data Repository backup to the same host](#) (see page 42).
3. [Configure Data Repository](#) (see page 44).

More information:

[Install the Data Repository Component](#) (see page 32)

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 38)

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 38)

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 38)

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).

Chapter 5: Installing the Data Aggregator Component

This section contains the following topics:

[Data Aggregator System Requirements](#) (see page 49)

[How to Prepare for a Data Aggregator Installation](#) (see page 52)

[Install Data Aggregator from the Command Line](#) (see page 56)

Data Aggregator System Requirements

Virtual and SAN Environments

Review the document titled *CA Support Statement for Running CA Performance Management Products in Virtualization and SAN Environments* on CA Support Online:

https://support.ca.com/phpdocs/7/5637/CA_Support_Statement_for_Running_CA_NV_M_Products_in_Virtualization_Environments.pdf

The document discusses the CA policies for installing and operating Performance Management products on virtualized servers or Storage Array Networks (SAN). You must be logged in as a registered user to view documentation on CA Support Online.

Small Deployments

A small deployment supports 15,000 metrics per second, for monitoring up to 100,000 devices and components. This deployment can be configured on the following four-host virtual machine environment:

- One CA Performance Center host
- One Data Repository host
- One Data Aggregator host
- One Data Collector host

The following information describes the minimum system requirements for installing the Data Aggregator host:

Component	Requirement
Operating system	Red Hat Enterprise Linux 5.x for x64 Red Hat Enterprise Linux 6.x for x64 Note: Upgrading an existing RH5 Vertica installation to an RH6 installation is not supported.
Processor	2 dedicated virtual CPUs
Processor speed	2.0 GHz
Memory (RAM)	16 GB
Disk space	50 GB Note: If Data Export will be used for one hour of data export, a second, separate 5 GB partition is required.
Disk input/output	100 Mega Bytes/second

Note: The data loading process has been rearchitected to persist all polled data temporarily to files on Data Aggregator before loading the data into Data Repository. Polled data files are removed once they are loaded into Data Repository. As a result, the amount of available disk space on Data Aggregator will fluctuate over time.

Medium Deployments

A medium deployment supports 75,000 metrics per second, for monitoring up to 500,000 monitored devices and components. This deployment can be configured on the following six-host physical machine environment:

- One CA Performance Center host
- Three Data Repository hosts in a cluster
- One Data Aggregator host
- One Data Collector host

The following information describes the minimum system requirements for installing the Data Aggregator host:

Component	Requirement
Operating system	Red Hat Enterprise Linux 5.x for x64 Red Hat Enterprise Linux 6.x for x64 Note: Upgrading an existing RH5 Vertica installation to an RH6 installation is not supported.
Processor	8 cores
Processor speed	2.0 GHz
Memory (RAM)	32 GB
Disk space	50 GB Note: If Data Export will be used for one hour of data export, a second, separate 25 GB partition is required.
Disk input/output	100 Mega Bytes/second

Note: The data loading process has been rearchitected to persist all polled data temporarily to files on Data Aggregator before loading the data into Data Repository. Polled data files are removed once they are loaded into Data Repository. As a result, the amount of available disk space on Data Aggregator will fluctuate over time.

Large Deployments

A large deployment supports 150,000 metrics per second, for monitoring up to one million monitored devices and components. This deployment can be configured on the following seven-host physical machine environment:

- One CA Performance Center host
- Three Data Repository hosts in a cluster
- One Data Aggregator host
- Two Data Collector hosts

The following information describes the minimum system requirements for installing the Data Aggregator host:

Component	Requirement
Operating system	Red Hat Enterprise Linux 5.x for x64 Red Hat Enterprise Linux 6.x for x64 Note: Upgrading an existing RH5 Vertica installation to an RH6 installation is not supported.
Processor	12 cores
Processor speed	2.5 GHz
Memory (RAM)	48 GB
Disk space	100 GB Note: If Data Export will be used for one hour of data export, a second, separate 5 GB partition is required.
Disk input/output	100 Mega Bytes/second

Note: The data loading process has been rearchitected to persist all polled data temporarily to files on Data Aggregator before loading the data into Data Repository. Polled data files are removed once they are loaded into Data Repository. As a result, the amount of available disk space on Data Aggregator will fluctuate over time.

How to Prepare for a Data Aggregator Installation

Meet the following prerequisites before installing Data Aggregator:

1. Verify that Data Repository is set up and running.

2. Open both port numbers 8581 and 61616 on the Data Aggregator system. Remote access is required to this port.

Note: You can change the 61616 port number to another port after you install Data Aggregator.

3. Verify that Security Enhanced Linux (SELinux) is disabled on the computer where you are going to install Data Aggregator. By default, some Linux distributions have this feature enabled, which does not allow Data Aggregator to function properly. Disable SELinux or create a policy to exclude Data Aggregator processes from SELinux restrictions.

Note: For information about configuring an SELinux security policy, see the Red Hat documentation.

4. Verify that the directory where you are going to install has write privileges for your Data Aggregator user.
5. [Ensure that the user that is installing Data Aggregator has a ulimit value of at least 65536.](#) (see page 53)
6. (Optional) [Configure the sudo user account](#) (see page 54).
7. [Configure UTF-8 support](#) (see page 55).
8. [Review the installation considerations](#) (see page 56).

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 nfile 65536
# Added by Data Aggregator
* hard nfile 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 Prepare for a Data Aggregator Installation](#) (see page 52)

Configure the Sudo User Account for Data Aggregator (Optional)

Before you install Data Aggregator, log in as the root user. However, in some environments, unrestricted root user access is not available. If root user access is not available, a sudo user with access to a limited set of commands can install and run the software.

Follow these steps:

1. Log in to the computer where you want to install Data Aggregator as the root user.
2. Add the following command alias to the command alias section of the `/etc/sudoers` file:

```
Cmnd_Alias CA_DATAAGG = /tmp/installDA.bin, /etc/init.d/dadaemon,  
/opt/IMDataAggregator/uninstall
```

```
## Allows the Data Aggregator user to manage the Data Aggregator
```

```
dasudouser_name ALL = CA_DATAAGG
```

This command alias details the commands that the sudo user must be able to run.

The sudo user account is configured.

More Information:

[How to Prepare for a Data Aggregator Installation](#) (see page 52)

Configure UTF-8 Support

Configure the computer where you will install the component to support UTF-8 encoding. If UTF-8 encoding is not enabled, characters may not display properly during the installation.

The appropriate language packs are also required to support localized deployments.

Note: Some scripts that are used in the installation of selected components are not localized and run in English. For more information, see the *Localization Status Readme* file.

Follow these steps:

1. Do one of the following steps:

- a. Type the following command from a Korn or bash shell:

```
export LANG=LANG_value ; export LC_ALL=$LANG
```

LANG_value

Indicates the value of the language you want the product to support. The following variables are supported:

English: en_US.utf8

French: fr_FR.utf8

Japanese: ja_JP.utf8

Simplified Chinese: zh_CN.utf8

Traditional Chinese: zh_TW.utf8

For example:

```
export LANG=zh_TW.utf8 ; export LC_ALL=$LANG
```

- b. Type the following command from a Bourne shell:

```
LANG=LANG_value ; export LANG
```

```
LC_ALL=LANG_value ; export LC_ALL
```

For example:

```
LANG=zh_CN ; export LANG
```

```
LC_ALL=zh_CN ; export LC_ALL
```

The language variable is set.

More Information:

[How to Prepare for a Data Aggregator Installation](#) (see page 52)

Data Aggregator Considerations

Consider the following information before you install Data Aggregator:

- Data Aggregator can communicate with only one Data Repository at a time.
- In a multi-tenant deployment, Data Aggregator *can* be shared between tenants. The information for each tenant is secure and other tenants cannot view this information.

Install Data Aggregator from the Command Line

If you launch the Data Aggregator installer from a secure shell, no GUI environment installation is available. You can install Data Aggregator using command-line prompts (console mode) only. The installer automatically switches to console mode even if you specify GUI mode.

Note the following information when using response files:

- Both versions of response file can be used in GUI mode installation using the `-f` option.
- Response files cannot be used in the console mode (command line) installation at all. The installer ignores it.
- Only the GUI version response file can be used for silent mode 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 a local directory. For the sake of installation instructions, we assume that the installer is placed in the `/tmp` directory, however, this location is not required.
3. Change permissions for the installation file by typing the following command:

```
chmod a+x installDA.bin
```
4. Type the following command to access the `/tmp` folder:

```
cd /tmp
```
5. To run the console installation, do one of the following steps:
 - To run the installation as the root user, type the following command:

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

```
sudo ./installDA.bin -i console
```

The License Agreement opens.

6. Read the license agreement, accept the agreement, and click Next.
7. 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.
8. Enter an installation directory when prompted.

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 following 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 name or the IP address of any one of the three hosts that are participating in the cluster. The installer automatically determines the name and IP address of the remaining nodes.

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 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 the Data Repository users that you created:

New User Example	Password Example	Operating System User Account?	Vertica Database User Account?
dauser	dapass	No	Yes
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 console installer displays a message asking 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 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
- Data Aggregator is installed.
13. Verify that Data Aggregator is installed:
- To verify that the installation was successful, review the information in the CA_Infrastructure_Management_Data_Aggregator_Install_*timestamp*.log file. This log file is located in the Logs subdirectory of the directory where you installed Data Aggregator.

- Verify that the Data Aggregator service is started and running. Open your web browser on a computer where you have HTTP access to Data Aggregator. Navigate to the following address:

`http://data_aggregator:port/rest`

data_aggregator:port

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

The return is a list of hyperlinks for available web services. When you click a link, such as Monitoring Profiles, the XML content describing the selection displays.

- Open a command prompt and type the following command to verify that the ActiveMQ broker is running:

`/etc/init.d/activemq status`

14. (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*.

15. 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.

Chapter 6: Installing the Data Collector Component

This section contains the following topics:

[Data Collector System Requirements](#) (see page 61)

[How to Prepare for a Data Collector Installation](#) (see page 64)

[Install Data Collector from the Command Line](#) (see page 68)

Data Collector System Requirements

Virtual and SAN Environments

Review the document titled *CA Support Statement for Running CA Performance Management Products in Virtualization and SAN Environments* on CA Support Online:

https://support.ca.com/phpdocs/7/5637/CA_Support_Statement_for_Running_CA_NV_M_Products_in_Virtualization_Environments.pdf

The document discusses the CA policies for installing and operating Performance Management products on virtualized servers or Storage Array Networks (SAN). You must be logged in as a registered user to view documentation on CA Support Online.

Small Deployments

A small deployment supports 15,000 metrics per second, for monitoring up to 100,000 devices and components. This deployment can be configured on the following four-host virtual machine environment:

- One CA Performance Center host
- One Data Repository host
- One Data Aggregator host
- One Data Collector host

The following information describes the minimum system requirements for installing the Data Collector host:

Component	Requirement
Operating system	Red Hat Enterprise Linux 5.x for x64 Red Hat Enterprise Linux 6.x for x64 Note: Upgrading an existing RH5 Vertica installation to an RH6 installation is not supported.
Processor	2 dedicated virtual CPUs
Processor speed	2.0 GHz
Memory (RAM)	12 GB
Disk space	50 GB
Disk input/output	100 Mega Bytes/second

Medium Deployments

A medium deployment supports 75,000 metrics per second, for monitoring up to 500,000 monitored devices and components. This deployment can be configured on the following six-host physical machine environment:

- One CA Performance Center host
- Three Data Repository hosts in a cluster
- One Data Aggregator host
- One Data Collector host

Note: These requirements are for a Data Collector instance that is monitoring up to 500,000 devices and components. In practice, multiple lighter Data Collector instances can be substituted, as long as the combined capacity does not exceed 500,000 monitored devices and components. For example, you can use five virtual Data Collector instances to approximate one large one. You can see how many devices and components Data Collector is monitoring in the Data Collector List view. For more information about the Data Collector List view, see the *Data Aggregator Administrator Guide*.

The following information describes the minimum system requirements for installing the Data Collector host:

Component	Requirement
Operating system	Red Hat Enterprise Linux 5.x for x64 Red Hat Enterprise Linux 6.x for x64 Note: Upgrading an existing RH5 Vertica installation to an RH6 installation is not supported.
Processor	4 cores
Processor speed	2.0 GHz
Memory (RAM)	32 GB
Disk space	50 GB
Disk input/output	100 Mega Bytes/second

Large Deployments

A large deployment supports 150,000 metrics per second, for monitoring up to one million monitored devices and components. This deployment can be configured on the following seven-host physical machine environment:

- One CA Performance Center host
- Three Data Repository hosts in a cluster
- One Data Aggregator host
- Two Data Collector hosts

Note: This information is specific to *each* of the two Data Collector instances. In practice, multiple lighter Data Collector instances can be substituted, as long as the combined capacity does not exceed one million monitored devices and components. For example, you can use five virtual Data Collector instances to approximate one large one. You can see how many devices and components Data Collector is monitoring in the Data Collector List view. For more information about the Data Collector List view, see the *Data Aggregator Administrator Guide*.

The following information describes the minimum system requirements for installing the Data Collector host:

Component	Requirement
Operating system	Red Hat Enterprise Linux 5.x for x64 Red Hat Enterprise Linux 6.x for x64 Note: Upgrading an existing RH5 Vertica installation to an RH6 installation is not supported.
Processor	4 cores
Processor speed	2.5 GHz
Memory (RAM)	32 GB
Disk space	50 GB
Disk input/output	100 Mega Bytes/second

How to Prepare for a Data Collector Installation

Meet the following prerequisites before you install Data Collector:

1. Verify that port number 61616 is open on the Data Aggregator system. This port lets Data Collector communicate with Data Aggregator.
2. Verify that Security Enhanced Linux (SELinux) is disabled on the computer where you are going to install Data Collector. By default, some Linux distributions have this feature enabled, which does not allow Data Collector to function properly. Disable SELinux or create a policy to exclude Data Collector processes from SELinux restrictions.

Note: For information about configuring an SELinux security policy, see the Red Hat documentation.

3. (Optional) [Configure the sudo user account](#) (see page 65).
4. Ensure that your desired tenant and corresponding IP domain are provisioned in CA Performance Center. While a single IP domain can be associated with more than one Data Collector, each Data Collector can have only one IP domain assigned to it.

Note: If you are not deploying multi-tenancy, use the Default Tenant and the Default Domain. For more information about creating tenants and IP domains, see the *CA Performance Center Administrator Guide*.

5. Configure UTF-8 support.
6. [Set a unique hostname for each host](#) (see page 66).
7. [Review the installation considerations](#) (see page 67).

Configure the Sudo User Account for Data Collector (Optional)

Before you install Data Collector, log in as the root user. However, in some environments, unrestricted root user access is not available. If you cannot log in as the root user, a sudo user with access to a limited set of commands can install and run the software.

Follow these steps:

1. Log in to the computer where you want to install Data Collector as the root user.
2. Add the following command alias to the command alias section of the `/etc/sudoers` file:

```
Cmnd_Alias CA_DATACOLL = /tmp/install.bin, /etc/init.d/dcmd,  
/opt/IMDataCollector/Uninstall/Uninstall
```

```
## Allows the Data Collector user to manage the Data Collector
```

```
sudouser_name ALL = CA_DATACOLL
```

This command alias details the commands that the sudo user must be able to run.

The sudo user account is configured.

More information:

[How to Prepare for a Data Collector Installation](#) (see page 64)

Configure UTF-8 Support

Configure the computer where you will install the component to support UTF-8 encoding. If UTF-8 encoding is not enabled, characters may not display properly during the installation.

The appropriate language packs are also required to support localized deployments.

Note: Some scripts that are used in the installation of selected components are not localized and run in English. For more information, see the *Localization Status Readme* file.

Follow these steps:

1. Do one of the following steps:

- a. Type the following command from a Korn or bash shell:

```
export LANG=LANG_value ; export LC_ALL=$LANG
```

LANG_value

Indicates the value of the language you want the product to support. The following variables are supported:

English: en_US.utf8

French: fr_FR.utf8

Japanese: ja_JP.utf8

Simplified Chinese: zh_CN.utf8

Traditional Chinese: zh_TW.utf8

For example:

```
export LANG=zh_TW.utf8 ; export LC_ALL=$LANG
```

- b. Type the following command from a Bourne shell:

```
LANG=LANG_value ; export LANG
```

```
LC_ALL=LANG_value ; export LC_ALL
```

For example:

```
LANG=zh_CN ; export LANG
```

```
LC_ALL=zh_CN ; export LC_ALL
```

The language variable is set.

More Information:

[How to Prepare for a Data Aggregator Installation](#) (see page 52)

Set a Unique Hostname for the Data Collector Host

Set a unique hostname for the computer where you plan to install Data Collector.

Follow these steps:

1. As the root user, log in to the computer where you are going to install Data Collector and verify the unique hostname on the computer.

The hostname for the computer must be associated with the IP address and *not* the loopback address of 127.0.0.1.

2. Verify that the following lines appear in the `/etc/hosts` file on the computer:

```
Do not remove the following line, or various programs
# that require network functionality will fail.
127.0.0.1 localhost.localdomain localhost
::1 localhost6.localdomain6 localhost6
IP address of your host YourHostName YourHostName.ca.com
```

3. If the hostname required any changes, type the following command after you make the changes:

```
service network restart
```

The `/etc/hosts` file is configured correctly.

The unique host name is set.

More information:

[How to Prepare for a Data Collector Installation](#) (see page 64)

Data Collector Considerations

Consider the following information before you install Data Collector:

- Install Data Collector on a separate computer from the server where Data Aggregator is installed.
- You can install more than one Data Collector. However, each Data Collector must be installed on a separate computer.
- Data Collector can support only one Data Aggregator.
- In a multi-tenant deployment, Data Collector is not shared among tenants. However, a tenant can have more than one Data Collector.
- In a multi-tenant environment where a managed service provider is monitoring devices for multiple tenants, you can take the following steps:
 - Install Data Collector at the MSP site.
Note: This setup requires Data Collector to gain access through a tenant firewall to poll the devices that are being managed.
 - Install Data Collector at each tenant site.

- If Data Aggregator is IPv6 only, Data Collector must support the IPv6 protocol.

To verify that Data Collector supports IPv6, take the following steps:

- On the Data Aggregator host, type the following command to find the IPv6 address of the computer:

```
> ifconfig
```
- On the Data Collector host, type the following command to ensure that Data Collector can contact Data Aggregator using its IPv6 address:

```
> ping6 ipv6_address_of_Data_Aggregator
```

Install Data Collector from the Command Line

If you launch the Data Collector installer from a secure shell or console and you are not running an X Window System on the computer where you want to install Data Collector, the installer launches in console mode (command line) automatically. Install Data Collector after you install Data Aggregator.

Note the following information when using response files:

- Both versions of response file can be used in GUI mode installation using the `-f` option.
- Response files cannot be used in the console mode (command line) installation at all. The installer ignores it.
- Only the GUI version response file can be used for silent mode installation.

If you are installing more than one Data Collector, install each Data Collector instance on a separate computer.

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 where Data Aggregator is installed *and* you are running an X Window System, 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 have HTTP access to where Data Aggregator is installed and you are *not* running an X Window System, open a command prompt from the computer where you want to install Data Collector. Type the following command to download the installation package to the /tmp directory:

`wget /tmp -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.

- If you do *not* have HTTP access to 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.

3. Transfer the install.bin file to a local directory on the computer where you want to install Data Collector. For the sake of installation instructions, we assume that the installer is placed in the /tmp directory, however, this location is not required.

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

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

```
cd /tmp
```

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

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

6. To run the console installation, do one of the following steps:

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

```
./install.bin -i console
```

- To run the installation as the sudo user, type the following command

```
sudo ./install.bin -i console
```

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.
11. When the installer prompts you for the Data Aggregator host information, enter either the IP address or the hostname for Data Aggregator that will be associated with Data Collector.

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, started, and connects to Data Aggregator.

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

13. Review the `/opt/IMDataCollector/Logs/CA_Infrastructure_Management_Data_Collector_install_timestamp.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 doing the following actions:
 - a. Log in to CA Performance Center as the global administrator.
 - 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. 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 another tenant and an IP domain for this Data Collector in the Assign Data Collector dialog and click Save.Data Collector is installed. Data Collector is running as the installation owner.
16. Type the following command to verify that Data Collector is running:

service dcmd status

Chapter 7: Installing CA Performance Center

This section contains the following topics:

[Deployment Options](#) (see page 73)

[Installation Considerations](#) (see page 76)

[Install CA Performance Center](#) (see page 84)

Deployment Options

Any CA Performance Center deployment requires additional servers to host data sources. Data source system requirements vary widely. Consult the data source Release Notes for deployment guidelines and system requirements.

CA has performed comprehensive testing of three different deployment scenarios with different numbers of managed items. The number of managed items translates to a rate of metric data that is processed:

- 15,000 metrics per second; 100,000 managed items (smaller deployment)
- 75,000 metrics per second; 500,000 managed items (medium-size deployment)
- 150,000 metrics per second; 1,000,000 managed items (large deployment)

The system requirements in this document reflect the differing requirements of these three environments.

System Requirements

Prepare a dedicated server for the CA Performance Center installation. Be sure that the server meets the requirements for your deployment size.

For All Deployments

The software can be installed in any filesystem to which the root user has write access. Typically, it is installed in /opt.

Verify that Security Enhanced Linux (SELinux) is disabled on the server where you plan to install CA Performance Center. By default, some Linux distributions have this feature enabled, which does not allow the product to function properly.

By default, the MySQL database is installed to /opt/CA/MySQL/, but you can select another location. Make sure you have enough disk space allocated to the selected filesystem to support a database.

The minimum supported screen resolution is 1280x1024.

DNS resolution and time synchronization are required. For more information, see the *Installation Guide*.

Adobe® Flash® must be installed on any client that accesses the product user interface. You can download it from the Adobe website.

A web browser is required. The following browser applications are supported:

- Microsoft Internet Explorer versions 9, 10, 11
- Mozilla Firefox (current version)
- Google Chrome (current version)

Virtual and SAN Environments

Review the document titled *CA Support Statement for Running CA Performance Management Products in Virtualization and SAN Environments* on CA Support Online:

https://support.ca.com/phpdocs/7/5637/CA_Support_Statement_for_Running_CA_NV_M_Products_in_Virtualization_Environments.pdf

The document discusses the CA policies for installing and operating Performance Management products on virtualized servers or Storage Array Networks (SAN). You must be logged in as a registered user to view documentation on CA Support Online.

15,000 Metrics per Second (Smaller Deployment)

Setting or Component	Description
Operating system	Red Hat Enterprise Linux 5.x for x64 Red Hat Enterprise Linux 6.x for x64
Processor	4 dedicated virtual CPUs
Processor Speed	2.0 GHz
Memory (RAM)	16 GB
Disk Space	56 GB
Disk Input/Output	100 Mega Bytes/second

75,000 Metrics per Second (Medium-Size Deployment)

Setting or Component	Description
Operating System	Red Hat Enterprise Linux 5.x for x64 Red Hat Enterprise Linux 6.x for x64
Processor	12 cores
Processor Speed	2.0 GHz
Memory (RAM)	32 GB
Disk Space	130 GB
Disk Input/Output	100 Mega Bytes/second

150,000 Metrics per Second (Large Deployment)

Setting or Component	Description
Operating System	Red Hat Enterprise Linux 5.x for x64 Red Hat Enterprise Linux 6.x for x64
Processor	24 cores
Processor Speed	2.5 GHz
Memory	32 GB
Disk Space	223 GB
Disk Input/Output	100 MegaBytes/second

Installation Considerations

Consider the following factors before you install CA Performance Center:

- Server prerequisites are carefully detailed in the Release Notes.
- CA Performance Center installation is required to deploy the Event Manager.
The Event Manager is installed and configured automatically as part of CA Performance Center installation.
- When installing on Linux, administrator-level access is required. If you do not have root access to the server, the user account you are using must be sudo-enabled. For more information, see [Linux User Account Requirements](#) (see page 78).
- Verify that Security Enhanced Linux (SELinux) is disabled on the computer where you plan to install CA Performance Center. By default, some Linux distributions have this feature enabled, which does not allow the product to function properly. Disable SELinux, or create a policy to exclude CA Performance Center daemons from SELinux restrictions.

Note: For information about configuring an SELinux security policy, see the Red Hat documentation.

- The installation package does not include antivirus software. We recommend installing your preferred antivirus software to protect your networking environment.

Important! To avoid database corruption, exclude the installation directory, and all its subdirectories, from antivirus scans. Prevent scanning by a local instance of an antivirus client and scanning by a remote antivirus instance.

- By default, the installation directory on Linux is /opt/CA/PerformanceCenter. The Setup program lets you select another location.
- For CA Performance Center to work properly in a firewall-protected environment, consider the communication ports that must be open. Configure firewalls to open the ports that are required for CA Performance Center and for any data sources that you plan to register.

Consult the *Installation Guide* of each data source for the list of required ports.

- CA Performance Center requires DNS resolution. If DNS is not configured, add system entries to the /etc/hosts file on your server manually.
- Time synchronization using NTP is also required. Start the NTP daemon on Linux if it is not running. For more information, see [Verify Time Synchronization](#) (see page 80).

CA Performance Center Communication Ports

CA Performance Center uses multiple ports to communicate with various components, particularly data sources. In addition, some of the products and components that integrate with CA Performance Center have specific port requirements.

Important! For any firewall that protects this server, open the required ports and protocols for the data sources you are deploying. The product documentation for each data source provides a list of required ports and protocols.

Each data source uses unique ports. However, the following communication ports must be open to allow communications between CA Performance Center and various products or components:

TCP/HTTP 80

Enables synchronization with CA Network Flow Analysis to retrieve configuration data.

TCP 3306

Enables communications from the MySQL database (inbound) on the console.

TCP/HTTP 8181

Enables communications between client computers and the CA Performance Center server. Enables console communications with data sources.

TCP/HTTP 8281

Enables communications between the Event Manager, which is installed automatically with the CA Performance Center software, and the data sources.

TCP/HTTP 8381

Enables communications between client computers and the CA Performance Center server. Also enables login using the Single Sign-On authentication component.

TCP 8481

Enables communications between the Device Manager and Console services.


TCP/HTTP 8681

Enables synchronization with CA Network Flow Analysis to retrieve device data.

TCP/HTTP 8581

Enables synchronization with CA Performance Management.

The following diagram illustrates basic port usage:



Be sure to consult the *Administrator Guide* or *Installation Guide* of the data sources you install for their port requirements.

Linux User Account Requirements

Administrative privileges are required to install the software. A user logged in as 'root' typically installs the software on the Linux server. However, in some environments, unrestricted root user access is not available.

If you cannot log in as root, you can use a user account with 'sudo' enabled to install and run the software. This user account must be granted the ability to use sudo to run a required, limited set of commands as the root user.

You can set up several sudo user accounts on the Linux server. These "super-user" accounts are necessary for starting and stopping CA Performance Center daemons.

The following command alias details the commands that the sudo user must be able to run. Run this command alias to set up your `/etc/sudoers` file:

```
Cmdn_Alias CA_PERFCENTER = /tmp/CAPerfCenterSetup.bin,  
/etc/init.d/caperfcenter_console, /etc/init.d/caperfcenter_devicemanager,  
/etc/init.d/caperfcenter_eventmanager, /etc/init.d/caperfcenter_sso,  
/etc/init.d/mysql,  
/opt/CA/PerformanceCenter/Tools/bin/npcshell.sh,  
/opt/CA/PerformanceCenter/SsoConfig,  
/opt/CA/PerformanceCenter/Uninstall_MySql,  
/opt/CA/PerformanceCenter/Uninstall_PerformanceCenter,  
/opt/CA/PerformanceCenter/Uninstall_SS0  
sudouser ALL = CA_PERFCENTER
```

You can add a reference to this alias for your sudo user in the `/etc/sudoers` file. You must then use the following command to install:

```
sudo location of CAPerfCenterSetup.bin
```

Increase Thread Allocation in Large Deployments

The System Requirements in the *Release Notes* advise you to provision a server with multiple CPUs in larger-size deployments. We recommend changing the `thread_concurrency` parameter to account for the number of CPUs on the server.

Follow these steps:

1. Log in to the server where you have installed CA Performance Center.
2. Edit the following file:

`/etc/my.cnf`
3. Search for the 'thread_concurrency' parameter.
4. Change the number of threads to equal two times the total number of CPUs on the server.
5. Save the file.
6. Stop and restart the mysql daemon by entering the following commands:

```
service mysql stop  
service mysql start
```

Verify Time Synchronization

Time synchronization using the network time protocol (NTP) daemon is required for CA Performance Center and is recommended for all data source consoles. On Linux servers, the NTP daemon ensures that the clocks on the hosts are synchronized for timing purposes. Verify that the daemon is running on the CA Performance Center host server.

Follow these steps:

1. Open a console and type the following command:

```
$ chkconfig --list ntpd
```

If the NTP daemon is installed, the output resembles the following example:

```
ntp 0:off 1:off 2:on 3:on 4:off 5:on 6:off
```

The output indicates the runlevels where the daemon runs.

2. Verify that the current runlevel of the system (usually 3 or 5) has the NTP daemon set to 'on'. If you do not know the current runlevel, type the following commands to find it:

```
$ runlevel  
N 3
```

If the current runlevel does not have the NTP daemon enabled, enable it by typing the following command:

```
$ chkconfig ntpd on
```

3. Type the following command to start the NTP daemon manually:

```
$ /etc/init.d/ntpd start
```

The daemon is started.

Modify Maximum Memory Usage for Each Service

Modify the maximum memory usage for the CA Performance Center daemons to enable them to run effectively. For system requirements, see the *Release Notes*.

You can configure memory allocation during or after the installation.

Follow these steps:

1. Log in to the server where you have installed CA Performance Center.
2. Enter the following command:

```
more /proc/meminfo
```

The total memory usage of each process is displayed.

3. Make a note of the total memory.
4. Modify the maximum memory for a daemon as follows:
 - a. Edit the following file:

*/Installation Directory/PerformanceCenter/Service
Subdirectory/conf/wrapper.conf*

Note: The Service subdirectory is one of the following options:

- PC (Console daemon)
 - DM (Device Manager daemon)
 - EM (Event Manager daemon)
- b. Search for the parameter 'wrapper.java.maxmemory'.
 - c. Change the current value. For example, in a small deployment, set it to '3072' (the units are MB).
 - d. Save the file.
 - e. Stop and restart each daemon by entering the following commands:

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

Note: The service name is one of the following options:

- caperfcenter_console
- caperfcenter_devicemanager
- caperfcenter_eventmanager

The maximum amount of memory is configured for a deployment of your scalability requirements.

More information:

[Install Data Collector from the Command Line](#) (see page 68)

[Install Data Aggregator from the Command Line](#) (see page 56)

Configure UTF-8 Support

Configure the computer where you will install the component to support UTF-8 encoding. If UTF-8 encoding is not enabled, characters may not display properly during the installation.

The appropriate language packs are also required to support localized deployments.

Note: Some scripts that are used in the installation of selected components are not localized and run in English. For more information, see the *Localization Status Readme* file.

Follow these steps:

1. Do one of the following steps:

- a. Type the following command from a Korn or bash shell:

```
export LANG=LANG_value ; export LC_ALL=$LANG
```

LANG_value

Indicates the value of the language you want the product to support. The following variables are supported:

English: en_US.utf8

French: fr_FR.utf8

Japanese: ja_JP.utf8

Simplified Chinese: zh_CN.utf8

Traditional Chinese: zh_TW.utf8

For example:

```
export LANG=zh_TW.utf8 ; export LC_ALL=$LANG
```

- b. Type the following command from a Bourne shell:

```
LANG=LANG_value ; export LANG
```

```
LC_ALL=LANG_value ; export LC_ALL
```

For example:

```
LANG=zh_CN ; export LANG
```

```
LC_ALL=zh_CN ; export LC_ALL
```

The language variable is set.

More Information:

[How to Prepare for a Data Aggregator Installation](#) (see page 52)

Third-Party Software

Except for anti-virus, system management, and time-synchronization software, do not install third-party software, especially third-party network monitoring software, on the same server as CA Performance Center. Third-party software can interfere with the monitoring abilities of the CA system and could void the warranty.

If you install third-party software on a CA system, CA Support might ask you to uninstall this software before troubleshooting an issue on the server.

Set the Limit on the Number of Open Files on CA Performance Center

Verify that the user account that is installing CA Performance Center 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 CA Performance Center. 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 CA Performance Center and add the following lines:

```
# Added by Performance Center
* soft nofile 65536
# Added by Performance Center
* hard nofile 65536
```

Note: Restart CA Performance Center for these changes to take affect. If you are upgrading, the upgrade process automatically restarts CA Performance Center.

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

```
ulimit -n
```

The command returns the ulimit number that you specified earlier.

Install CA Performance Center

Use the CA Performance Center Setup program to install and configure the database and website. The Setup program provides a graphical user interface, but you can also install from the command line.

Follow these steps:

1. Log in to the server as root, or use the sudo account you have configured for the installation.

If the sudo account is not configured properly, you are prompted to set it up with the appropriate permissions. For more information, see [Linux User Account Requirements](#) (see page 78).

2. Change to the directory where you extracted the contents of the installation package.

```
cd/tmp
```

3. Change the directory permissions.

```
chmod +x /tmp/CAPerfCenterSetup.bin
```

4. Run the following command:

```
./CAPerfCenterSetup.bin -i console
```

The installation begins in Console mode.

5. Select your preferred language, and press Enter.

The Welcome message advises you to close any programs that are running on the server. If necessary, enter 'Quit' to exit the setup program and close programs.

6. Press Enter to continue with the installation.

The License Agreement is displayed. Press Enter until you are prompted to accept it.

7. Enter Y to accept the license agreement.

You are prompted to 'Select Installation Directory'.

The default installation directory is /opt/CA.

8. Enter an absolute path to a new installation directory, or press Enter to accept the default.

You are prompted to 'Set the maximum memory allocation for the CA Performance Center services.'

9. Set the size, in megabytes, of the memory allocation for the Console, Device Manager, and Event Manager services, and press Enter.

For more information about memory allocation, see [Modify Maximum Memory Usage for Each Service](#) (see page 80).

You are prompted to 'Select a Location for the MySQL Data Directory'.

10. Type a location for the data directory that the database uses, or press Enter to accept the default. The default is /opt/CA/MySQL/data.

Important! Verify that the drive you select has sufficient space available. 40 GB of space are required for the database.

11. Press Enter to continue.

You are prompted to 'Select a Location for the MySQL Temp Directory'.

12. Type a location for the tmp directory that the database uses, or press Enter to accept the default. The default is /opt/CA/MySQL/tmp.

You are prompted to start the installation.

13. Press Enter to continue.

You are prompted to 'Review Installation Settings'. The settings that you have selected for the installation are displayed. The information also includes guidance for the amount of disk space required.

14. Review the selected settings, and enter 'back' if you want to change any settings. You can enter 'back' multiple times to back up multiple steps.

15. Press Enter to begin the installation.

Note: Enter 'quit' to exit the setup program without installing the software.

Text appears to indicate that the installation is progressing.

16. The following Linux daemons are created and started during the installation:

capformancecenter_console

Is the console daemon. Uses port 8181.

capformancecenter_devicemanager

Is the Device Manager daemon. Uses port 8481.

capformancecenter_eventmanager

Is the Event Manager daemon. Uses port 8281.

caperformancecenter_sso

Is the Single Sign-On daemon. Uses port 8381.

mysql

Is the database daemon. Uses port 3306.

When the installation has completed, a message states that the program has been installed successfully.'

17. Press Enter to exit the installation program.

Install Support for Non-English Languages

CA Performance Center and its data sources provide support for multiple languages. The administrator can select a preferred language for each unique product operator. Language packs take advantage of operating system support for localized environments.

However, product operators with a language preference other than English might not be able to view dashboard data in reports by default. You might need to install additional fonts on the server as a separate step.

Follow the standard instructions for installing fonts on your operating system. CA Performance Center reporting and export options are already available for the following fonts:

- Arial
- Arial Unicode MS
- Liberation Sans
- Sans
- Meiryo UI
- AR PL ShanHeiSun Uni
- SimSun
- Sazanami Mincho
- AR PL ZenKai Uni
- Baekmuk Batang

To use a different font on your operating system, perform the following procedure.

Follow these steps:

1. Follow the steps in the *Installation Guide* to install the CA Performance Center software.
2. Bring up the Linux package manager by executing the following command:

```
pirut
```
3. Select Languages from the left side of the list on the Browse tab.
4. Select the language(s) to install from the right side of the list.
5. Click Apply.

Chapter 8: Post-Installation Configuration Options

This section contains the following topics:

[How to Complete the Installation](#) (see page 89)

How to Complete the Installation

Perform the following optional and recommended steps after you install Data Aggregator:

1. (Optional) [Set up autostart on Data Repository](#) (see page 89).
2. (Recommended) [Configure the automatic recovery of the Data Aggregator process.](#) (see page 93)
3. (Optional) [Modify the maximum memory usage for Data Aggregator and Data Collector components after installation.](#) (see page 94)
4. (Optional) [Modify the external ActiveMQ memory limit after installation](#) (see page 97).
5. (Optional) [Change the opened port number on the Data Aggregator host.](#) (see page 98)

Set Up Autostart on Data Repository (Optional)

You can set up autostart on Data Repository. If autostart is set up and you reboot the computer where Data Repository is installed, Data Repository starts automatically.

Important! This feature may not work properly if Data Repository did not shut down gracefully. If the database did not shut down gracefully, the database might require manual intervention during startup to restore the last good epoch. If the Vertica database does not start automatically after an ungraceful shutdown, use admintools to start it manually.

Data Aggregator stops automatically when Data Repository becomes inaccessible. Restart Data Aggregator manually once Data Repository is online again. To restart Data Aggregator, open a command prompt and type the following command:

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

Follow these steps:

1. To become the Linux user account for the database administrator user, type the following command:

```
su - dradmin
```
2. To navigate to the
`/opt/vertica/config/users/Linux_user_account_for_database_administrator_user`
directory, type the following command:

```
cd  
/opt/vertica/config/users/Linux_user_account_for_database_administrator_user
```
3. To copy the `installed.dat` file to a new file called `dbinfo.dat`, type the following command:

```
cp -p installed.dat dbinfo.dat
```

The `dbinfo.dat` file is created.
4. Be sure that the Linux user account for the database administrator user is set up with a passwordless ssh key:
 - 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. However, if you are asked if you want to continue connecting, enter Yes.
 - b. If you *are* prompted for a password, ignore the prompt and hit Ctrl+C.
5. To set up the Linux user account for the database administrator user with a passwordless ssh key, do the following steps:
 - a. To become the Linux user account for the database administrator user, type the following command:

```
su - dradmin_username
```
 - b. To generate a public key, type the following command. In a cluster installation, type this command on each host that is participating in the cluster:

```
ssh-keygen -N "" -t rsa -f ~/.ssh/id_rsa
```
 - c. Copy the contents of the public key to the `authorized_keys2` file on the same computer. In a cluster installation, copy the contents of the public key to the `authorized_keys2` file on each host that is participating in the cluster:

```
cat ~/.ssh/id_rsa.pub >> ~/.ssh/authorized_keys2
```

- d. (Cluster installation only) Copy the contents of the public key from each host to each of the other hosts:

- As the database administrator user on the first host, type the following command and copy the content of the file:

```
vi ~/.ssh/id_rsa.pub
```

- As the database administrator user on the second host, type the following command:

```
vi ~/.ssh/authorized_keys2
```

Paste the contents from the id_rsa.pub file on the first host to the end of the authorized_keys2 file on the second host.

- As the database administrator user on the third host, type the following command:

```
vi ~/.ssh/authorized_keys2
```

Paste the contents from the id_rsa.pub file on the first host to the end of the authorized_keys2 file on the third host.

To enable you to ssh from one host to another without being prompted for a password, repeat these steps for all hosts in the cluster.

- e. To set permissions for the authorized_keys2 file, type the following command. In a cluster environment, type these commands on each host in the cluster:

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

- f. As the root user, type the following commands to restart the ssh daemon. In a cluster environment, type this command on each host in the cluster:

```
su - root
/etc/init.d/sshd restart
```

- g. (Single-node installation only) To confirm that you are not prompted for a password, type the following commands:

```
su - dradmin
ssh dradmin@hostname ls /tmp
```

- h. (Cluster installation only) To confirm that you are not prompted for a password, type the following commands on the first host in the cluster:

```
su - dradmin
ssh dradmin@host1 ls /tmp
ssh dradmin@host2 ls /tmp
ssh dradmin@host3 ls /tmp
```

Repeat this step on each host in the cluster.

Important! If you do not set up the passwordless ssh key, you cannot configure autostart on Data Repository.

6. Type the following command:

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

The Administration Tools dialog opens.

7. Select (6) Configuration Menu and select OK.

8. Select (4) Set Restart Policy and select OK.

The Select Database dialog opens.

9. Select the database name and select OK.

The Select policy dialog opens.

10. Select 'always' when doing a single-node Data Repository installation. Select 'ksafe' when doing a cluster installation.

Select OK.

Note: In a single-node installation, 'always' means that, Data Repository automatically restarts when the system restarts. In a cluster installation, 'ksafe' means that, upon the system restarting, the Data Repository node automatically restarts if the database still has a status of 'UP'.

The Restart Policy setting is saved.

11. Select OK to close the Select policy dialog.

12. Return to the (M) Main Menu.

13. Select (E) Exit.

14. (Optional) Test that Data Repository starts when you reboot the computer where Data Repository is installed:

- a. Reboot the computer where Data Repository is installed.

Note: Log in as the root user or sudo user to reboot the computer.

- b. Become the Linux user account for the database administrator user. Type the following command:

```
su - dradmin
```

- c. Type the following command:

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

The Administration Tools dialog opens.

- d. Select (1) View Database Cluster State and select OK.

- e. The state is "UP." Select OK.

Note: Data Repository can take several minutes to start up after you reboot.

More information:

[How to Complete the Installation](#) (see page 89)

Configure the Automatic Recovery of the Data Aggregator Process (Recommended)

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.

You can configure the automatic recovery of the Data Aggregator process.

When Data Aggregator shuts down, an audit message is logged in the *Data Aggregator installation directory*/apache-karaf-2.3.0/shutdown.log file. When Data Aggregator becomes unavailable, Data Collector continues polling and Data Collector 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.

We recommend that you disable this cron job before you upgrade Data Aggregator. If you shut down Data Aggregator manually with the service dadaemon stop command, the cron job does not restart Data Aggregator automatically. Maintenance can be performed without having the cron job disrupt the system when it is expected to be down.

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.

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. If there are no cron jobs for the database administrator user, an empty file opens. Otherwise, the file contains existing cron job definitions.

3. Add the following lines to the file for the cron job:

```
EXECUTED_BY_CRON=1
* * * * * service dadaemon start > /dev/null
```

This line means that cron is going to issue a start command to Data Aggregator every minute.

If Data Aggregator is running, the start command is ignored.

If Data Aggregator is *not* running, the start command starts Data Aggregator. After it starts, Data Aggregator checks for the availability of Data Repository. If Data Repository is unavailable, Data Aggregator shuts down. This process is repeated until all Data Repository connectivity problems are resolved.

More information:

[How to Complete the Installation](#) (see page 89)

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: 33554432KB

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.

More information:

[How to Complete the Installation](#) (see page 89)

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> and <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.

More information:

[How to Complete the Installation](#) (see page 89)

Change the Opened Port Number on the Data Aggregator Host (Optional)

After you install Data Aggregator, you can change the port that is opened on the Data Aggregator host.

Note: You opened port 61616 before you installed Data Aggregator and Data Collector.

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.

2. Open a command prompt and type the following command to stop Data Aggregator:

```
service dadaemon stop
```

3. Type the following commands to remove the data directory and the local-jms-broker.xml file from the deploy directory:

```
rm -rf Data Aggregator installation directory/apache-karaf-2.3.0/data
rm -rf Data Aggregator installation
directory/apache-karaf-2.3.0/deploy/local-jms-broker.xml
```

4. Edit the local-jms-broker.xml file in the *Data Aggregator installation directory*/apache-karaf-2.3.0/jms directory:

- a. Locate the following lines:

```
<!-- The transport connectors ActiveMQ will listen to -->
<transportConnectors>
  <transportConnector name="openwire" uri="tcp://dahostname:61616"/>
```

- b. Replace 61616 with the port that you want to use for incoming connections on Data Aggregator.

5. Open a command prompt and type the following command to start Data Aggregator:

```
service dadaemon start
```

6. Wait a few minutes, then type the following command to verify that the port change is successful:

```
netstat -a | grep port
```

port

Is the port number that you specified previously for incoming connections on Data Aggregator.

7. If the port change is successful, Data Aggregator waits for incoming connections on that port. If Data Aggregator is not waiting for incoming connections, type the following command to review the karaf.log file for errors:

```
grep ERROR karaf.log
```

8. Resolve the errors.

9. 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*.

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

```
service dcmd stop
```

11. Type the following commands to remove the data directory and the local-jms-broker.xml file from the deploy directory:

```
rm -rf Data Aggregator installation directory/apache-karaf-2.3.0/data
rm -rf Data Aggregator installation
directory/apache-karaf-2.3.0/deploy/local-jms-broker.xml
```

12. Edit the local-jms-broker.xml file in the *Data Collector installation directory*/apache-karaf-2.3.0/jms directory:

- a. Locate the following lines:

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

- b. Replace 61616 with the port that you specified previously in the local-jms-broker.xml file on the Data Aggregator host.

13. Open a command prompt and type the following command to start Data Collector:

```
service dcmd start
```

14. Wait a few minutes, then type the following command to verify that the port change is successful:

```
netstat -a | grep port
```

port

Is the port number that you specified in a previous step for incoming connections on Data Aggregator.

If the port change is successful, you should see a connection between Data Aggregator and Data Collector. If you do not see a connection, type the following command to review the karaf.log file for errors:

```
grep ERROR karaf.log
```

15. Resolve the errors.

The opened port number on the Data Aggregator host is changed.

More information:

[How to Complete the Installation](#) (see page 89)

Chapter 9: Troubleshooting

This section contains the following topics:

[Troubleshooting: Data Aggregator Synchronization Failure](#) (see page 101)

[Troubleshooting: CA Performance Center Cannot Contact Data Aggregator](#) (see page 102)

[Troubleshooting: Data Collector Installs But Does Not Appear in the Data Collector List Menu](#) (see page 103)

[Troubleshooting: Vertica Fails to Install in a Cluster Environment](#) (see page 104)

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:
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.

Troubleshooting: Vertica Fails to Install in a Cluster Environment

Symptom:

Vertica fails to install in my cluster environment.

Solution:

Set up passwordless SSH for the Vertica Linux database administrator user and then retry the installation. Do the following steps to set up passwordless SSH:

1. Open a console and log in to the Data Repository host as the Vertica Linux database administrator user.

2. 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
```

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

```
ssh-copy-id -i database_admin_user@remotehost
```

remotehost

Is another host in the cluster where you are trying to copy the SSH ID.

4. To verify that passwordless ssh is set up correctly, login to the remote host from the local host:

```
ssh database_admin_user@remotehost ls
```

5. Repeat steps 1-4 for each pair of hosts.

Note: A three-node cluster requires six variations of the previous steps.

If the passwordless SSH has been set up successfully, you are not prompted for a password. You also see a directory listing from the 'ls command'.