

HP Vertica Analytics Platform 6.1.x

Installation Guide

Doc Revision 3

Copyright© 2006-2013 Hewlett-Packard

Date of Publication: Monday, October 28, 2013



Contents

Installation Guide	5
Syntax Conventions	6
Installation Overview and Checklist	7
About HP Vertica-created Linux Users and their Privileges	9
Before You Install HP Vertica	12
Platform Requirements and Recommendations	12
Configuring Hardware and Installing Linux.....	14
Install Linux.....	14
Allocate Swap Space	14
Disable CPU Frequency Scaling	15
Configure the I/O Scheduler	15
Prepare Disk Storage Locations	17
Verify sudo and Allow Passwordless SSH Access for dbadmin	20
Verify pam_limits.so.....	21
Update tzdata Package and Set the Default Time Zone	21
Set the Host Locale (language).....	24
Verify That NTP daemon is Running.....	25
Remove nonessential applications	26
Configuring the Network	26
Define the Loopback Address	26
Configure Host Name Resolution	27
Check Subnet Masks.....	29
Disable Firewalls	30
Optionally Run Spread on Separate Control Network	30
Provide root SSH Access to the Cluster.....	30
Ensure Ports are Available	30
Ensure That /dev/pts is Mounted.....	31
Installing HP Vertica	32
Installing HP Vertica Using the Command Line	33
Back Up Existing Databases.....	33
Download and Install the HP Vertica Install Package	34
Run the Install Script.....	35
About the install_vertica Script.....	37
Installing HP Vertica Silently	48
Installing HP Vertica on Amazon Web Services (AWS).....	51
Creating an HP Vertica Cluster Using MC	51
Prepare the Hosts	51
Create a Private Key File	53

Use MC's Cluster Installation Wizard	54
Validate Hosts and Create the Cluster.....	55
Create a Database on a Cluster.....	56
Installing and Configuring Management Console	58
Before You Install MC.....	58
Installing MC	59
Configuring MC	60
After You Install HP Vertica	63
Install the License Key	63
Verify Kernel and User Parameters Were Set	63
Optionally Install vsql Client Application on Non-cluster Hosts.....	65
Install the HP Vertica-Ganglia Monitoring Package	65
Resolve SUSE Spread Configuration Issues.....	65
Changing the IP Addresses of an HP Vertica Cluster	66
Install HP Vertica Documentation.....	69
Installing Client Drivers	71
Extend Your Installation Using HP Vertica Packages	71
Upgrading HP Vertica	72
Upgrading HP Vertica to a New Version	73
Upgrading Your HP Vertica License	76
Upgrading MC	76
Uninstalling HP Vertica	77
Uninstalling HP Vertica	78
Uninstalling MC	78
Troubleshooting the HP Vertica Install	80
Validation Scripts	80
vcupperf.....	80
vioperf.....	81
vnetperf.....	83

Installation Guide

Enable Port 5433	86
Enable Secure Shell (SSH) Logins	86
Increase Free Memory	87
Increase Maximum Number of Files Open	88
Increase max_map_count Parameter	89
Set Up Time Synchronization	90
Manage Large Grouped ROS Containers	91

Appendix: Time Zones 93

Using Time Zones With HP Vertica	93
Africa	95
America	97
Antarctica.....	103
Asia	103
Atlantic	107
Australia	107
Etc/GMT	108
Europe	109
Indian	111
Pacific	112

Copyright Notice 115

Installation Guide

Welcome to the HP Vertica Installation Guide. Read this guide to learn how to prepare for and install the HP Vertica server. This guide also provides instructions for installing the HP Vertica Management Console.

For information about installing client drivers, see the Vertica Programmer's guide.

Prerequisites

- This document assumes that you have become familiar with the concepts discussed in the Concepts Guide.
- To perform the procedures described in this document, you must have root password or sudo access (for all commands) for all nodes in your cluster.

Syntax Conventions

The following are the syntax conventions used in the HP Vertica documentation.

Syntax Convention	Description
Text without brackets/braces	Indicates content you type, as shown.
< <i>Text inside angle brackets</i> >	Represents a placeholder for which you must supply a value. The variable is usually shown in italics. See <i>Placeholders</i> below.
[Text inside brackets]	Indicates optional items; for example, CREATE TABLE <i>[schema_name.]table_name</i> The brackets indicate that the <i>schema_name</i> is optional. Do not type the square brackets.
{ Text inside braces }	Indicates a set of options from which you choose one; for example: QUOTES { ON OFF } indicates that exactly one of ON or OFF must be provided. You do not type the braces: QUOTES ON
Backslash \	Represents a continuation character used to indicate text that is too long to fit on a single line.
Ellipses ...	Indicate a repetition of the previous parameter. For example, option[, ...] means that you can enter multiple, comma-separated options. Showing ellipses in code examples might also mean that part of the text has been omitted for readability, such as in multi-row result sets.
Indentation	Is an attempt to maximize readability; SQL is a free-form language.
<i>Placeholders</i>	Represent items that must be replaced with appropriate identifiers or expressions and are usually shown in italics.
Vertical bar	Is a separator for mutually exclusive items. For example: [ASC DESC] Choose one or neither. You do not type the square brackets.

Installation Overview and Checklist

This page provides an overview of installation tasks.

Important notes

- HP Vertica supports only one running database per cluster.
- HP Vertica supports installation on one, two, or multiple nodes. The steps in *Installing HP Vertica* (page [33](#)) are the same, no matter how many nodes are in the cluster.
- Prerequisites listed in *Before You Install HP Vertica* (page [12](#)) are required for all HP Vertica configurations.
- Only one instance of HP Vertica can be running on a host at any time.
- To run the `install_vertica` script, you must be logged in as root, or sudo as a user with all privileges. You must run the script for all installations, including upgrades and single-node installations.

Carefully review and follow the instructions in all sections in this topic.

Installation scenarios

The four main scenarios for installing HP Vertica on hosts are:

- A single node install, where HP Vertica is installed on a single host. This form of install cannot be expanded to more hosts later on and is typically used for development or evaluation purposes.
- Installing to a cluster of physical host hardware. This is the most common scenario when deploying HP Vertica in a testing or production environment.
- Installing on Amazon Web Services (AWS). When you choose the recommended Amazon Machine Image (AMI), Vertica is installed when you create your instances. For more information, see the guide, *Using HP Vertica on Amazon Web Services*.

Note: Follow the instructions and steps in *Using HP Vertica on Amazon Web Services* for an AWS configuration rather than the steps for installation and upgrade that appear in this guide.

- Installing to a local cluster of virtual host hardware. Also similar to installing on physical hosts, but with network configuration differences.

Before you install

Before You Install HP Vertica (page [12](#)) describes how to construct a hardware platform and prepare Linux for HP Vertica installation. These preliminary steps are broken into two categories: Configuring Hardware and Installing Linux and Configuring the Network.

- Configuring Hardware and Installing Linux. This section describes how to set up hardware and software for one or more hosts for an HP Vertica database. The steps in this section are the same, no matter what type of install (single node or multiple physical node) you are performing.
 - Install Linux
 - Allocate swap space
 - Configure the I/O scheduler for catalog and data drives

- Prepare disk storage locations
- Verify sudo
- Set up host locale and default time zone
- Ensure NTP is running on startup
- Remove nonessential applications
- Configuring the Network. Once you have configured the operating system, you need to configure the network settings.
 - Define the loopback address
 - Configure hostname resolution
 - Check subnet masks
 - Disable firewalls
 - [Optional] Run spread on a separate network
 - Provide root and SSH access to the cluster
 - Ensure that port 5433 is not in use and that `/dev/pts` is mounted.

Install or upgrade HP Vertica

Once you have completed the steps in the ***Before You Install HP Vertica*** (page [12](#)) section, you are ready to run the install script.

Installing HP Vertica (page [33](#)) describes how to:

- Back up any existing databases.
- Download and install the HP Vertica RPM package.
- Install a cluster using the `install_vertica` script.
- [Optional] ***Create a properties file*** (page [48](#)) that lets you install HP Vertica silently.

Note: This guide provides additional ***manual procedures*** (page [80](#)) in case you encounter installation problems.

Upgrading HP Vertica to a New Version (page [73](#)) describes the steps for upgrading to a more recent version of the software.

Upgrading HP Vertica from Community Edition to Enterprise Edition (page [76](#)) describes the steps for upgrading HP Vertica to an evaluation or Enterprise Edition version of the software.

Post-installation tasks

After You Install HP Vertica (page [63](#)) describes subsequent steps to take after you've run the installation script. Some of the steps can be skipped based on your needs:

- Install the license key
- Verify that kernel and user parameters are correctly set
- Install the vsql client application on non-cluster hosts
- Install the HP Vertica-Ganglia monitoring package
- Resolve any SUSE10 issues during spread configuration.

- Use the HP Vertica documentation online, or download and install HP Vertica documentation. Find the online documentation and documentation packages to download at <http://www.vertica.com/documentation> (<http://www.vertica.com/documentation>).
- Install client drivers.
- Extend your installation with HP Vertica packages.
- **Install** (page [58](#)) or **upgrade** (page [76](#)) the Management Console.

Get started!

- Read the Concepts Guide for a high-level overview of the HP Vertica Analytics Platform.
- Proceed to the Tutorial: Setting up an Example Database in the Getting Started Guide, where you will be guided through setting up a database, loading sample data, and running sample queries.

About HP Vertica-created Linux Users and their Privileges

When you install HP Vertica, the installation script creates two Linux users and a Linux group:

- dbadmin—Administrative user
- spread—Spread daemon account
- verticadba—Group for DBA users

dbadmin and verticadba are the default names. If you want to change what these Linux accounts are called, you can do so using the installation script. You cannot change the name of the spread account. See *The install_vertica Script* (page [37](#)) for details.

This topic describes the Linux accounts that the installer creates and configures so HP Vertica can run.

Before you install HP Vertica

The user who runs the HP Vertica installer must have sudo privileges on all cluster nodes. See the following topics for more information:

- **Installation Overview and Checklist** (page [7](#))
- **Verify sudo and Allow Passwordless SSH Access for dbadmin** (page [20](#))

When you install HP Vertica

The dbadmin user. The Linux dbadmin user owns the database catalog and data storage on disk. When you run the install script, HP Vertica creates this user on each node in the database cluster, adds dbadmin to the Linux dbadmin and verticadba groups, and configures the account as follows:

- Configures and authorizes dbadmin for passwordless SSH between all cluster nodes. SSH must be installed and configured to allow passwordless logins. See **Enable Secure Shell (SSH) Logins** (page [86](#)).
- Sets the dbadmin user's BASH shell to `/bin/bash`, required to run scripts, such as `install_vertica` and the Administration Tools.

- Provides read-write-execute permissions on the following directories:
 - `/opt/vertica/*`
 - `/home/dbadmin`, the default directory for database data and catalog files (configurable through the install script)

Note: The HP Vertica installation script also creates an HP Vertica database superuser named `dbadmin`. They share the same name, but they are not the same; one is a Linux user and the other is an HP Vertica user. See `DBADMIN` user in the Administrator's Guide for information about the database superuser

The spread user. Spread is a run-time daemon that provides an efficient, fault-tolerant subsystem for exchanging messages. During the installation process, the script creates:

- A spread system account in the `verticadba` group (set to `/nologin`)
- A spread service in `/etc/init.d/spreadd`

The installer configures spread to run on HP Vertica with the following permissions:

- Open non-privileged sockets with default port access at 4803 (UDP, TCP). Also uses UNIX domain socket in `/tmp/5433`.
- Write messages to the log file, which defaults to `/dev/null`.

Note: You can change the default location for log output to `/opt/vertica/log/spread_<host>.log` by using the `install_vertica` script `-w` argument.

- Run the spread daemon executable on each node in the database cluster. The daemon resides in `/opt/vertica/spread/sbin/spread` with `rwxr-xr-x` permissions.

The spread service starts at system boot time using `/etc/init.d/spreadd`, which automatically changes from root user permissions to spread user permissions before the spread executable launches. HP Vertica will not run if spread is not running.

The verticadba group. The installer creates a Linux group for DBA users, called `verticadba`, which provides control over filesystem access in the `/opt/vertica/` directories. HP Vertica moves existing Linux users into `verticadba` with permissions set to 775. This setting grants full privileges to the DBA 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 and are owned by the `dbadmin` user.

After you install HP Vertica

Root or sudo privileges are not required to start or run HP Vertica after the installation process completes.

The `dbadmin` user can log in and perform HP Vertica tasks, such as creating a database, installing/changing the license key, or installing drivers. If `dbadmin` wants database directories in a location that differs from the default, the root user (or a user with sudo privileges) must create the requested directories and change ownership to the `dbadmin` user.

HP Vertica prevents administration from users other than the `dbadmin` user (or the user name you specified during the installation process if not `dbadmin`). Only this user can run the Administration Tools.

See also

Installation Overview and Checklist (page [7](#))

Before You Install HP Vertica (page [12](#))

Platform Requirements and Recommendations (page [12](#))

Enable Secure Shell (SSH) Logins (page [86](#))

Verify sudo and Allow Passwordless SSH Access for dbadmin (page [20](#))

The install_vertica Script (page [37](#))

Before You Install HP Vertica

Complete all of the tasks in this section before you install HP Vertica. When you have completed this section, proceed to ***Installing HP Vertica*** (page [33](#)).

Platform Requirements and Recommendations

You must verify that your servers meet the platform requirements described in the Supported Platforms document available at <http://www.vertica.com/documentation> (<http://www.vertica.com/documentation>). The Supported Platforms document details supported versions for the following:

- OS
- Browser
- HP Vertica driver compatibility
- Perl
- Python
- R
- Hadoop
- Various plug-ins

Hardware Recommendations

The HP Vertica Analytic Database is based on a massively parallel processing (MPP), shared-nothing architecture, in which the query processing workload is divided among all nodes of the Vertica database. HP highly recommends using a homogeneous hardware configuration for your HP Vertica cluster; that is, each node of the cluster should be similar in CPU, clock speed, number of cores, memory, and operating system version.

Note that HP has not tested HP Vertica on clusters made up of nodes with disparate hardware specifications. While it is expected that an HP Vertica database would functionally work in a mixed hardware configuration, performance will most certainly be limited to that of the slowest node in the cluster.

Detailed hardware recommendations are available in the ***HP Vertica Hardware Planning Guide*** http://my.vertica.com/docs/6.1.x/HardwareDocs/HP_Vertica%20Hardware%20Planning%20Guide.pdf.

BASH shell

All shell scripts included in HP Vertica must run under the BASH shell. If you are on a Debian system, then the default shell may be DASH. DASH is not supported. Change the shell for root and for the dbadmin user to BASH with the `chsh` command. For example:

```
# echo $SHELL
/bin/dash
# chsh
Changing shell for root.
```

```
New shell [/bin/dash]: /bin/bash
Shell changed.
```

Then, as root, change the symbolic link for `/bin/sh` from `/bin/dash` to `/bin/bash`:

```
# rm /bin/sh
# ls -l /bin/bash /bin/sh
```

Log out and back in for the change to take effect.

Filesystem Options

- HP Vertica supports the ext3 and ext4 file systems.
- Linux Logical Volume Manager (LVM) is **not supported**.
- The recommended disk block size is 4096 bytes.

General Platform Recommendations

- ext4 is recommended over ext3 for performance reasons.
- Use 2GB of swap space regardless of the amount of installed RAM.
- Place the database `/catalog` directory on the same drive as the OS.

Install the Latest Vendor Specific System Software

Install the latest vendor drivers for your hardware. For HP Servers, update to the latest versions for:

- HP ProLiant Smart Array Controller Driver (cciss)
- Smart Array Controller Firmware
- HP Array Configuration Utility (HP ACU CLI)

Data Storage Recommendations

- All internal drives connect to a single RAID controller.
- The RAID array should form one hardware RAID device as a contiguous `/data` volume.
- The I/O Scheduler should be configured for deadline.

Validation Utilities

Starting in Release 6.0, there are several validation utilities that you can use to validate the performance on prospective hosts. The utilities are installed when you install the HP Vertica RPM, but they can be used before you run the `install_vertica` script. See **Validation Scripts** for more details on running the utilities and verifying that your hosts meet the recommended requirements.

Oracle Enterprise Linux Notes

If you are using Oracle Enterprise Linux, then note the following:

- When installing Oracle Enterprise Linux, select the vanilla kernel. HP Vertica does not support the unbreakable kernel (kernels with a `ue/` suffix).
- You cannot disable SELinux at install time when installing Oracle Enterprise Linux. Disable SELinux after install:
 1. Edit `/etc/selinux/config`.

2. Change the line `SELINUX=enforcing` to `SELINUX=disabled`
3. Reboot.

SUSE and Debian Notes

Make sure the Linux command-line calculation tool, `bc`, is installed and in the user's path on all nodes in the cluster (commonly in `/usr/bin/bc`) by running one of the following commands:

- On SUSE and Redhat: `rpm -qi bc`
- On Debian: `dpkg -l bc`

Run `which bc` to verify that `bc` is in the user's path, as in the following command:

```
[dbadmin@myhost01 ~]$ which bc
/usr/bin/bc
```

Configuring Hardware and Installing Linux

The first set of steps you need to take involve setting up your hardware platform and installing and configuring Linux. These steps are essentially the same, no matter if you are installing to a single node, a cluster, or a cloud.

Install Linux

The same version of Linux must be installed on all host machines within a cluster. See the HP Vertica Supported Platforms document (available on the Docs tab of the **myVertica portal** <http://my.vertica.com/>) for details on what platforms are supported.

Important! Deploy HP Vertica as the only active process on each host—other than Linux processes or software explicitly approved by HP Vertica. HP Vertica cannot be collocated with other software.

Allocate Swap Space

HP Vertica recommends that you allocate 2GB of swap space regardless of the amount of installed RAM. Larger swap space is acceptable, although unnecessary.

Note: Do not place a swap file on a disk containing the HP Vertica data files. If a host has only two disks (boot and data), put the swap file on the boot disk.

Disable CPU Frequency Scaling

CPU frequency scaling is a hardware and software feature that helps computers conserve energy by slowing the processor when the system load is low, and speeding it up again when the system load increases. This feature can impact system performance, since raising the CPU frequency in response to higher system load does not occur instantly. Always disable this feature on the database hosts to prevent it from interfering with HP Vertica's performance.

You disable CPU scaling in your host's system BIOS. There may be multiple settings in your host's BIOS that you need to adjust in order to completely disable CPU frequency scaling. Consult your host hardware's documentation for details of entering the system BIOS and disabling CPU frequency scaling.

Disabling CPU scaling in the kernel

If you cannot disable CPU scaling through the system BIOS, you can limit the impact of CPU scaling by disabling the scaling through the Linux kernel or setting the CPU frequency governor to always run the CPU at full speed.

Warning: This method is not reliable, as some hardware platforms may ignore the kernel settings. **The only reliable method is to disable CPU scaling in BIOS.**

The method you use to disable frequency depends on the CPU scaling method being used in the Linux kernel. See your Linux distribution's documentation for instructions on disabling scaling in the kernel or changing the CPU governor.

Configure the I/O Scheduler

The Linux kernel can use several different I/O schedulers to prioritize disk input and output. Most Linux distributions use the Completely Fair Queuing (CFQ) scheme by default, which gives input and output requests equal priority. This scheduler is efficient on systems running multiple tasks that need equal access to I/O resources. However, it can create a bottleneck when used on the drives containing HP Vertica's catalog and data directories, since it gives write requests equal priority to read requests, and its per-process I/O queues can penalize processes making more requests than other processes.

Instead of the CFQ scheduler, configure your hosts to use either the Deadline or NOOP I/O scheduler for the drives containing the catalog and data directories:

- The Deadline scheduler gives priority to read requests over write requests. It also imposes a deadline on all requests. After reaching the deadline, such requests gain priority over all other requests. This scheduling methods helps prevent processes from becoming starved for I/O access. The Deadline scheduler is best used on physical media drives (disks using spinning platters), since it attempts to group requests for adjacent sectors on a disk, lowering the time the drive spends seeking.

- The NOOP scheduler uses a simple FIFO approach, placing all input and output requests into a single queue. This scheduler is best used on solid state drives (SSDs). Since SSDs do not have a physical read head, no performance penalty exists when accessing non-adjacent sectors.

Failure to use one of these schedulers for the drives containing HP Vertica's catalog and data directories can result in slower database performance. Other drives on the system (such as the drive containing swap space, log files, or the Linux system files) can still use the default CFQ scheduler (although you should always use the NOOP scheduler for SSDs).

There are two ways for you to set the scheduler used by your disk devices:

- 1 Write the name of the scheduler to a file in the `/sys` directory
--or--
- 2 Use a kernel boot parameter

Changing the Scheduler through the `/sys` Directory

You can view and change the scheduler Linux uses for I/O requests to a single drive using a virtual file under the `/sys` directory. The name of the file that controls the scheduler a block device uses is:

```
/sys/block/deviceName/queue/scheduler
```

Where *deviceName* is the name of the disk device, such as `sda` or `cciss\!c0d1` (the first disk on an HP RAID array). Viewing the contents of this file shows you all of the possible settings for the scheduler, with the currently-selected scheduler surrounded by square brackets:

```
# cat /sys/block/sda/queue/scheduler
noop deadline [cfq]
```

To change the scheduler, write the name of the scheduler you want the device to use to its scheduler file. You must have root privileges to write to this file. For example, to set the `sda` drive to use the deadline scheduler, run the following command as root:

```
# echo deadline > /sys/block/sda/queue/scheduler
# cat /sys/block/sda/queue/scheduler
noop [deadline] cfq
```

Changing the scheduler immediately affects the I/O requests for the device. The Linux kernel starts using the new scheduler for all of the drive's input and output requests.

Note: While tests have shown no problems caused by changing the scheduler settings while HP Vertica is running, you should strongly consider shutting down any running HP Vertica database before changing the I/O scheduler or making any other changes to the system configuration.

Changes to the I/O scheduler made through the `/sys` directory only last until the system is rebooted, so you need to add commands that change the I/O scheduler to a startup script (such as those stored in `/etc/init.d`). You also need to use a separate command for each drive on the system whose scheduler you want to change.

You may prefer to use this method of setting the I/O scheduler over using a boot parameter if your system has a mix of solid-state and physical media drives, or has many drives that do not store HP Vertica catalog and data directories.

Changing the Scheduler with a Boot Parameter

Use the `elevator` kernel boot parameter to change the default scheduler used by all disks on your system. This is the best method to use if most or all of the drives on your hosts are of the same type (physical media or SSD) and will contain catalog or data files. You can also use the boot parameter to change the default to the scheduler the majority of the drives on the system need, then use the `/sys` files to change individual drives to another I/O scheduler. The format of the elevator boot parameter is:

```
elevator=schedulerName
```

Where `schedulerName` is `deadline`, `noop`, or `cfq`. You set the boot parameter using your bootloader (`grub` or `grub2` on most recent Linux distributions). See your distribution's documentation for details on how to add a kernel boot parameter.

Prepare Disk Storage Locations

You must create and specify directories in which to store your catalog and data files (physical schema). You can specify these locations when you install or configure the database, or later during database operations.

The directory you specify for your catalog files (the catalog path) is used across all nodes. That is, if you specify `/home/catalog` for your catalog files, HP Vertica will use `/home/catalog` as the catalog path on all nodes. The catalog directory should always be separate from any data files.

Note: Do not use a shared directory for more than one node. Data and catalog directories must be distinct for each node. Multiple nodes must not be allowed to write to the same data or catalog directory.

The same is true for your data path. If you specify that your data should be stored in `/home/data`, HP Vertica ensures this is the data path used on all database nodes.

Do not use a single directory to contain both catalog and data files. You can store the catalog and data directories on different drives, which can be either on drives local to the host (recommended for the catalog directory) or on a shared storage location, such as an external disk enclosure or a SAN.

Both the catalog and data directories must be owned by the database administrator.

Before you specify a catalog or data path, be sure to create these locations on all nodes of your database.

Disk Space Requirements for HP Vertica

In addition to actual data stored in the database, HP Vertica requires disk space for several data reorganization operations, such as mergeout and managing nodes in the cluster. For best results, HP recommends that disk utilization per node be no more than sixty percent (60%) for a `K-Safe=1` database to allow such operations to proceed.

In addition, disk space is temporarily required by certain query execution operators, such as hash joins and sorts, in the case when they have to spill to disk. Such operators might be encountered during queries, recovery, refreshing projections, and so on. The amount of disk space needed in this manner (known as temp space) depends on the nature of the queries, amount of data on the node and number of concurrent users on the system. By default, any unused disk space on the data disk can be used as temp space. However, it is possible and recommended to provision temp space separate from data disk space. See ***Configuring Disk Usage to Optimize Performance*** (page [20](#)).

Disk Space Requirements for Management Console

You can install MC on any node in the cluster, so there are no special disk requirements for MC—other than disk space you would normally allocate for your database cluster. See ***Disk Space Requirements for HP Vertica*** (page [17](#)).

Specifying Disk Storage Location During Installation

There are three ways to specify the disk storage location. You can specify the location when you:

- Install HP Vertica
- Create a database using the Administration Tools
- Install and configure Management Console

To specify the disk storage location when you install:

When you install HP Vertica, the `data_directory` parameter in the `install_vertica` (page [37](#)) script lets you specify a directory to contain database data and catalog files. The script defaults to the Database Administrator's default home directory: `/home/dbadmin`.

You should *replace this default* with a directory that has adequate space to hold your data and catalog files.


Before you create a database, verify that the data and catalog directory exists on each node in the cluster. Also verify that the directory on each node is owned by the database administrator.

Notes

- Catalog and data path names must contain only alphanumeric characters and cannot have leading space characters. Failure to comply with these restrictions will result in database creation failure.
- HP Vertica refuses to overwrite a directory if it appears to be in use by another database. Therefore, if you created a database for evaluation purposes, dropped the database, and want to reuse the database name, make sure that the disk storage location previously used has been completely cleaned up. See *Creating and Configuring Storage Locations* for details.

Specifying Disk Storage Location During Database Creation

When you invoke the Create Database command in the Administration Tools, a dialog box allows you to specify the catalog and data locations. These locations must exist on each host in the cluster and must be owned by the database administrator.



When you click **OK**, HP Vertica automatically creates the following subdirectories:

```
catalog-pathname/database-name/node-name_catalog/
data-pathname/database-name/node-name_data/
```

For example, if you use the default value (the database administrator's home directory) of `/home/dbadmin` for the Stock Exchange example database, the catalog and data directories are created on each node in the cluster as follows:

```
/home/dbadmin/Stock_Schema/stock_schema_node1_host01_catalog
/home/dbadmin/Stock_Schema/stock_schema_node1_host01_data
```

Notes

- Catalog and data path names must contain only alphanumeric characters and cannot have leading space characters. Failure to comply with these restrictions will result in database creation failure.
- HP Vertica refuses to overwrite a directory if it appears to be in use by another database. Therefore, if you created a database for evaluation purposes, dropped the database, and want to reuse the database name, make sure that the disk storage location previously used has been completely cleaned up. See *Creating and Configuring Storage Locations* for details.

Specifying Disk Storage Location on MC

You can use the MC interface to specify where you want to store database metadata on the cluster in the following ways:

- When you configure MC the first time
- When you create new databases using on MC

See **Configuring Management Console** (page [60](#)).

Configuring Disk Usage to Optimize Performance

Once you have created your initial storage location, you can add additional storage locations to the database later. Not only does this provide additional space, it lets you control disk usage and increase I/O performance by isolating files that have different I/O or access patterns. For example, consider:

- Isolating execution engine temporary files from data files by creating a separate storage location for temp space.
- Creating labeled storage locations and storage policies, in which selected database objects are stored on different storage locations based on measured performance statistics or predicted access patterns.

See [Creating and configuring storage locations](#) for details.

Using Shared Storage With HP Vertica

If using shared SAN storage, ensure there is no contention among the nodes for disk space or bandwidth.

- Each host must have its own catalog and data locations. Hosts cannot share catalog or data locations.
- Configure the storage so that there is enough I/O bandwidth for each node to access the storage independently.

Verify sudo and Allow Passwordless SSH Access for dbadmin

Verify that sudo is available on all cluster hosts and that you have privileges to use it on these hosts by executing the following command at a shell prompt:

```
# which sudo
/usr/bin/sudo
```

If sudo is not installed, browse to the **Sudo Main Page** <http://www.gratisoft.us/sudo/> and install sudo on all hosts.

When you use sudo to install HP Vertica, the user that performs the installation must have privileges on all nodes in the cluster.

Allow root permission via sudo

Configuring sudo with privileges for the individual commands can be a tedious and error-prone process; thus, the HP Vertica documentation does not include every possible sudo command that you can include in the sudoers file. Instead, HP recommends that you temporarily elevate the sudo user to have all privileges for the duration of the install.

Note: See the sudoers and visudo man pages for the details on how to write/modify a sudoers file.

To allow root sudo access on all commands as any user on any machine, use visudo as root to edit the `/etc/sudoers` file and add this line:

```
## Allow root to run any commands anywhere
root    ALL=(ALL) ALL
```

After the installation completes, remove (or reset) sudo privileges to the pre-installation settings.

Allow passwordless SSH access for the dbadmin user

The dbadmin user must be authorized for passwordless ssh. In typical installs, you won't need to change anything; however, if you set up your system to disallow passwordless login, you'll need to enable it for the dbadmin user. See *Enable secure shell (SSH) logins* (page [86](#)).

Verify pam_limits.so

On some systems the pam module called `pam_limits.so` is not set in the file `/etc/pam.d/su`. When it is not set, it prevents the conveying of limits (such as open file descriptors) to any command started with `su -`

In particular, the HP Vertica init script would fail to start HP Vertica because it calls the Administration Tools to start a database with the `su -` command. This problem was first noticed on Debian systems, but the configuration could be missing on other Linux distributions.

To resolve this issue, modify the `install_vertica` script and append this line to the `/etc/pam.d/su` file:

```
session required pam_limits.so
```

See also

pam_limits http://man.he.net/man8/pam_limits man page.

Update tzdata Package and Set the Default Time Zone

Update tzdata Package

The tzdata package is a public-domain time zone database that is pre-installed on most linux systems. The tzdata package is updated periodically for time-zone changes across the world. HP recommends that you update to the latest tzdata package before installing or updating HP Vertica.

Update your tzdata package with the following command:

- For RedHat based systems: `yum update tzdata`
- For Debian and Ubuntu systems: `apt-get install tzdata`

Setting the Default Time Zone

When a client receives the result set of a SQL query, all rows contain data adjusted, if necessary, to the same time zone. That time zone is the default time zone of the initiator node unless the client explicitly overrides it using the SQL `SET TIME ZONE` command described in the SQL Reference Manual. The default time zone of any node is controlled by the `TZ` environment variable. If `TZ` is undefined, the operating system time zone.

If your operating system timezone is not set to the desired timezone of the database then make sure that the Linux environment variable `TZ` is set to the desired value on all cluster hosts. Typically, the `TZ` variable is the same on all cluster hosts but this is not required; you can configure a cluster to provide convenient client connections from multiple time zones as explained below.

The installer returns a warning if the `TZ` variable is not set. If your operating system timezone is appropriate for your database, then the operating system timezone is used and the warning can be safely ignored.

Configuring multiple time zones

As a convenience to clients in multiple time zones, you can designate specific nodes as having specific default time zones. This flexibility lets a client connect to a specific node and execute SQL statements using the default SQL time zone. In other words, there is no need for the client to use the SQL `SET TIME ZONE` command.

If you do this, however, be aware that:

- The `install_vertica` script (described in *Installing HP Vertica* (page [33](#))) issues warnings when the `TZ` (and `LANG`) environment variables are not the same on all cluster hosts.
- A client that does not use the SQL `SET TIME ZONE` command could get unexpected query results if it connects to the wrong node.
- The timestamps in the each node's log files are based on its default time zone and could be different across the cluster.

Setting the time zone on a host

Important: If you explicitly set the `TZ` environment variable at a command line before you start the Administration Tools, the current setting will not take effect. The Administration Tools uses SSH to start copies on the other nodes, so each time SSH is used, the `TZ` variable for the startup command is reset. `TZ` must be set in the `.profile` or `.bashrc` files on all nodes in the cluster to take affect properly.

You can set the time zone several different ways, depending on the Linux distribution or the system administrators' preferences.

- To set the system time zone on Red Hat and SUSE Linux systems, edit:
`/etc/sysconfig/clock`
- To set the `TZ` variable, edit `/etc/profile` or `/root/.bashrc` and add the following line:
`export TZ=time-zone-name`

The following time zone names are recognized by HP Vertica as valid settings for the SQL time zone (the `TIME ZONE` run-time parameter).

Note: The names listed here are for convenience only and could be out of date. Refer to the **Sources for Time Zone and Daylight Saving Time Data** <http://www.twinsun.com/tz/tz-link.htm> page for precise information.

These names are not the same as the names shown in `/opt/vertica/share/timezonesets`, which are recognized by HP Vertica in date/time input values. The TIME_ZONE names shown below imply a local daylight-savings time rule, where date/time input names represent a fixed offset from UTC.

In many cases there are several equivalent names for the same zone. These are listed on the same line. The table is primarily sorted by the name of the principal city of the zone.

In addition to the names listed in the table, HP Vertica accepts time zone names of the form *STDoffset* or *STDoffsetDST*, where *STD* is a zone abbreviation, *offset* is a numeric offset in hours west from UTC, and *DST* is an optional daylight-savings zone abbreviation, assumed to stand for one hour ahead of the given offset. For example, if `EST5EDT` were not already a recognized zone name, it would be accepted and would be functionally equivalent to USA East Coast time. When a daylight-savings zone name is present, it is assumed to be used according to USA time zone rules, so this feature is of limited use outside North America. Be wary that this provision can lead to silently accepting bogus input, since there is no check on the reasonableness of the zone abbreviations. For example, `SET TIME_ZONE TO FOOBANKO` works, leaving the system effectively using a rather peculiar abbreviation for GMT.

Time Zone
Africa (on page 95)
America (on page 97)
Antarctica (on page 103)
Asia (on page 103)
Atlantic (on page 107)
Australia (on page 107)
CET
EET
Etc/GMT (on page 108)
Europe (on page 109)
Factory
GMT GMT+0 GMT-0 GMT0 Greenwich Etc/GMT Etc/GMT+0 Etc/GMT-0 Etc/GMT0

Etc/Greenwich
Indian (on page 111)
MET
Pacific (on page 112)
UCT Etc/UCT
UTC Universal Zulu Etc/UTC Etc/Universal Etc/Zulu
WET

See Also

Using Time Zones with HP Vertica (page [93](#))

Set the Host Locale (language)

Each host has a system setting for the Linux environment variable `LANG`. `LANG` determines the locale category for native language, local customs, and coded character set in the absence of the `LC_ALL` and other `LC_` environment variables. `LANG` can be used by applications to determine which language to use for error messages and instructions, collating sequences, date formats, and so forth.

To change the `LANG` setting for the database administrator, edit `/home/dbadmin/.profile` on all cluster hosts and set the environment variable; for example:

```
export LANG=en_US.UTF-8
```

The `LANG` setting controls the following in HP Vertica:

- OS-level errors and warnings, for example, "file not found" during COPY operations.
- Some formatting functions, such as `TO_CHAR` and `TO_NUMBER`. See also Template Patterns for Numeric Formatting.

The `LANG` setting does not control the following:

- HP Vertica-specific error and warning messages. These are always in English at this time.
- Collation of results returned by SQL issued to HP Vertica. This must be done using a database parameter instead. See Implement Locales for International Data Sets section in the Administrator's Guide for details.

Note: If the `LC_ALL` environment variable is set, it supersedes the setting of `LANG`.

Verify That NTP daemon is Running

The network time protocol (NTP) daemon needs to be running on all of the hosts in the cluster to ensure their clocks are synchronized. The spread daemon relies on all of the nodes in the having their clocks synchronized for timing purposes. If your nodes do not have NTP running, the installation can fail with a spread configuration error as well as other potential errors.

Note: Different Linux distributions refer to the NTP daemon in different ways. For example, SUSE and Debian refer to it as `ntp`, while CentOS and Red Hat refer to it as `ntpd`. If the following commands produce an error, try using `ntp` in place of `ntpd`.

To check if your hosts are configured to run the NTP daemon on startup, run the following command:

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

Note: Debian does not install `chkconfig` by default, but does offer it as an optional package. See the Debian documentation for information on how to obtain and install packages.

If `chkconfig` command produces an error similar to `ntpd: unknown service`, check to ensure your Linux distribution does not refer to the NTP daemon as `ntp` rather than `ntpd`. If it does not, you need to install the NTP daemon package before you can configure it. Consult your Linux documentation for instructions on how to locate and install packages.

If the NTP daemon is installed, output should resemble the following:

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

The output indicates the runlevels where the daemon runs. 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, you can find it using the `runlevel` command:

```
$ runlevel
N 3
```

If the current runlevel does not have the NTP daemon enabled, you can enable it by running the command:

```
$ chkconfig ntpd on
```

This configures the NTP daemon to run in the current runlevel. You then need to either reboot the host, or manually start the NTP daemon before continuing the installation process. You can start the daemon manually using the command:

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

NTP on Management Console

HP Vertica assigns each request to the agent a self-signed certificate, which includes a timestamp. To increase security and protect against password replay attacks, the timestamp is valid for several seconds only, after which it expires and prevents access to MC by issuing an HTTP 401 ACCESS DENIED error.

To avoid being blocked out of MC, synchronize time on the the hosts in your HP Vertica cluster, including the host on which you installed MC (which may or may not be in the database cluster). To recover from loss or lack of synchronization, resync system time and the Network Time Protocol. See **Set Up Time Synchronization** (page [90](#)).

Your browser might display alerts about the self-signed certificate when you connect to MC using the IP address of one of the hosts. You can ignore the warnings and accept the certificate, or you can extract the self-signed certificate and replace it with a real signed certificate, which you must generate outside of the application. You then upload a new SSL certificate on the **MC Settings** page and Restarting MC.

See Also

Generating certifications and keys for MC

Remove nonessential applications

For optimal performance, HP Vertica is designed to use all available resources on each host machine. HP recommends that you:

- Remove or disable all non-essential applications from cluster hosts
- Deploy HP Vertica as the only active process on each host—other than Linux processes

Configuring the Network

The group of steps involve configuring the network. These steps will differ depending on your installation scenario. A single node installation requires little network configuration, since the single instance of the HP Vertica server does not need to communication with other nodes in a cluster. For cluster and cloud install scenarios, there are several decisions you need to make regarding your configuration.

Define the Loopback Address

Make sure that the `/etc/hosts` file exists and that it contains the loopback address 127.0.0.1 with only the name `localhost` in both fully qualified and unqualified forms. For example:

```
127.0.0.1          localhost.localdomain localhost
```

Notes

Unless you are installing HP Vertica on a single host machine with no plans to expand to a multiple-node cluster or import or export data to or from another database:

- Do not use the loopback address `127.0.0.1` or the name `localhost` in a node definition.
- Do not assign the cluster node host name to `127.0.0.1`.

See Also

Configure hostname resolution (page [27](#))

Configure Host Name Resolution

It is essential that host names of all hosts used to create a cluster resolve correctly. Invalid host name resolution is a common source of configuration problems. Your HP Vertica database will not install or work properly unless the host name resolution is correct.

Even single-node installations must use only static addresses or permanently-leased DHCP addresses. The only exception is if you use `localhost` for a single host machine, in which case you cannot expand your cluster later.

Setting up cluster hosts

This procedure sets up the minimal host name resolution required for an HP Vertica installation.

To set up each host machine that becomes part of the database cluster:

1 Set up the `/etc/hosts` file on each node in the cluster.

Make sure that the `/etc/hosts` file includes all of the hosts that become part of the cluster. For example, if the hosts are named `host01`, `host02`, `host03`, and `host04`, the `/etc/hosts` file on each host looks like this:

```
# Do not remove the following line, or various programs
# that require network functionality will fail.
127.0.0.1          localhost.localdomain localhost
192.168.13.128     host01
192.168.13.129     host02
192.168.13.130     host03
192.168.13.131     host04
```

Note: Replace the example IP addresses shown here with the real addresses. Do not remove the loopback address (`127.0.0.1`).

2 Optionally configure separate private and public interfaces.

HP Vertica supports machine configuration with multiple network interfaces. For example, you might want to use one as a private network interface for internal communication among cluster hosts (the ones supplied via the `-s` option to `install_vertica`) and a separate one for client connections. For simplicity, HP recommends that each interface resolve to a different hostname. However, if a host has multiple IP addresses that resolve to the same hostname, the IP address used for the private interface must be the one listed first in `/etc/hosts` on all nodes in the cluster.

For example, if hostname `host01` resolved to both `192.168.10.100` and `172.16.69.100` where the `192.168.10.100` address is a private network to be used for HP Vertica, then list the `192.168.10.100` address first in `/etc/hosts` file.

All private interfaces should be on the same subnet and have the same broadcast address.

It is not necessary to have hostnames for the private interfaces (e.g., you can use IP addresses without hostname resolution instead). Use the IP addresses for the hosts for the `-s` parameter of `install_vertica`: `-s 192.168.233.101,192.168.233.102,192.168.233.103`.

3 Verify that the hostname resolution works correctly.

Make sure that the host identifies itself correctly. For example, on `host01`:

```
$ /bin/hostname -f
localhost.localdomain
```

`localhost.localdomain` is incorrect. In this case, you must set the hostname manually.

For details, refer to the operating system instructions or consult with the system administrator.

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

The command `/bin/hostname host01` sets the hostname to `host01` until system reboot.

4 Edit the `/etc/sysconfig/network` file to set `HOSTNAME` to `host01`:

```
HOSTNAME=host01
```

- Ping the host. For example:

```
$ ping host01
PING host01 (192.168.13.128) 56(84) bytes of data.
64 bytes from host01 (192.168.13.128): icmp_seq=0 ttl=64 time=0.086
ms
:
```

- Compare the ping command output to the hostname. The output of the hostname command must be identical to that of the ping command.
- Check that `/etc/sysconfig/network` contains the same hostname.

Setting the `HOSTNAME` environment variable

Some Linux installations require that the `HOSTNAME` environment variable be set correctly.

On each host machine that becomes part of the database cluster:

- 1 If the `HOSTNAME` environment variable is not defined, add the following line to your `/etc/profile` or `/etc/bashrc` file:

```
export HOSTNAME=`hostname`
```

- 2 If the `HOSTNAME` environment variable is already defined, make sure that it matches the hostname output:

```
# echo $HOSTNAME
host01
```

If they do not match, change the value of `HOSTNAME` in `/etc/profile` or `/etc/bashrc`:

```
# export HOSTNAME=`/bin/hostname -f`
```

Important notes

- Network configuration is exactly the same for single nodes as for multi-node clusters, with one special exception. If you install HP Vertica on a single host machine that is to remain a permanent single-node configuration (such as for development or Proof of Concept), you can install HP Vertica using `localhost` as the value for `-s`. Do not use the hostname `localhost` in a node definition if you are likely to add nodes to the configuration later.

Note: Installation using `localhost` is the only configuration mode where a static IP address is not required. All other configurations require static IP addresses. However, it is not possible to expand the cluster from a `localhost` installation. You need to re-install using IP address or hostname.

- If you are using a host with multiple network interfaces, use the IP address assigned to the NIC that is connected to the other cluster hosts, not the NIC that is used for client connections.
- Use a dedicated gigabit switch. If you do not, a WARNING is returned during installation, and performance could be severely affected. The warning reads "Low throughput <host1> to <host2>: value Mbps; check network interface/switch configuration."
- Do not use DHCP dynamically-assigned IP addresses for the private network. Use only static addresses or permanently-leased DHCP addresses. If your private and public interfaces have different hostnames, then you might be able to use DHCP for the public interfaces; however, this configuration is not currently supported or recommended by HP.

Examples

The following is an example of an `/etc/hosts` configuration that is recommended. It uses different host names for private and public network interfaces. All hosts for private interface are next to each other followed by all hosts for the public interface.

```
127.0.0.1      localhost.localdomain localhost
192.168.0.1    privhost1.verticacorp.com qa1
192.168.0.2    privhost2.verticacorp.com qa2
192.168.0.3    privhost3.verticacorp.com qa3
192.168.0.4    privhost4.verticacorp.com qa4

10.10.50.123   host01.verticacorp.com qa01
10.10.50.124   host02.verticacorp.com qa02
10.10.50.125   host03.verticacorp.com qa03
10.10.50.126   host04.verticacorp.com qa04
```

See also

Change the IP addresses of an HP Vertica cluster (page [66](#))

Check Subnet Masks

All hosts used in an HP Vertica database must have the same subnet mask.

Disable Firewalls

Firewalls are not recommended for database hosts. If it is consistent with your security policy, you can disable the Linux firewall and SELinux. Follow the documentation for your specific OS distribution.

See also

Enable port 5433 (page [86](#))

Optionally Run Spread on Separate Control Network

If your query workloads are network intensive, you can use the `-s` parameter with the `install_vertica` (page [37](#)) script to allow spread communications to be configured on a subnet that is different from other HP Vertica data communications.

The `-s` parameter accepts either the `default` value or a broadcast network IP address (for example, `192.168.10.255`).

Provide root SSH Access to the Cluster

Verify that root can use Secure Shell (SSH) to log in (`ssh`) to all hosts that are included in the cluster. SSH (SSH client) is a program for logging into a remote machine and for running commands on a remote machine.

If you do not already have SSH installed on all hosts, log in as root on each host and install it now. You can download a free version of the SSH connectivity tools from **OpenSSH** (<http://www.openssh.com/>).

Ensure Ports are Available

Port 5433

HP Vertica requires that port 5433 be available for use for client connectivity (vsq, ODBC, JDBC, etc). Before you install HP Vertica, verify that port 5433 is not in use by running the following command as the root user:

```
netstat -atup
```

Note that this command will not list the port number if it is identified by a service defined in `/etc/services` (`cat /etc/services` to see port definitions for these services). Some common services that might use port 5433 include `pyrrho`, `nibbler`, and `postgresql`.

The following is example output of the `netstat` command listed above with an instance of vertica running:

```
tcp        0      0  *:pyrrho          *:*
```

LISTEN

3855/vertica

Firewall considerations

HP Vertica requires several ports to be open on the local network. HP Vertica does not recommend placing a firewall between nodes (all nodes should be behind a firewall), but if you must use a firewall between nodes, ensure the following ports are available:

Port	Protocol	Service	Notes
22	TCP	sshd	Required by Administration Tools and the Management Console Cluster Installation (page 54) wizard.
5433	TCP	HP Vertica	HP Vertica client (vsq, ODBC, JDBC, etc) port.
5433	UDP	HP Vertica	HP Vertica spread monitoring.
5434	TCP	HP Vertica	Intra-cluster communication. See Intra-cluster Communication below.
5444	TCP	HP Vertica Management Console	MC-to-node and node-to-node (agent) communications port. See Changing MC or Agent Ports.
5450	TCP	HP Vertica Management Console	Port used to connect to MC from a web browser and allows communication from nodes to the MC application/web server. See Connecting to Management Console.
4803	TCP	Spread	Client connections.
4803	UDP	Spread	Daemon to Daemon connections.
4804	UDP	Spread	Daemon to Daemon connections.
4805	UDP	Spread	Monitor to Daemon connection.

Intra-cluster communication

HP Vertica opens a port for intra-cluster communication, such as during a plan. It opens the HP Vertica client port +1 (5434 by default) for this purpose. If you change the default client port, then you must open the firewall for the next port as well. If the port +1 from the default client port is not available, then HP Vertica opens a random port for intra-cluster communication.

Ensure That /dev/pts is Mounted

Make sure that /dev/pts is mounted. Installing HP Vertica on a host that is missing the mount point /dev/pts could result in the following error when you create a database:

TIMEOUT ERROR: Could not login with SSH. Here is what SSH said:
Last login: Sat Dec 15 18:05:35 2007 from node01

Installing HP Vertica

There are different paths you can take when installing HP Vertica. You can:

- Install HP Vertica on one or more hosts using the command line, and not use the Management Console.
- Install the Management Console, and from the Management Console install HP Vertica on one or more hosts by using the Management Console cluster creation wizard.
- Install HP Vertica on one or more hosts using the command line, then install the Management Console and import the cluster to be managed.

Installing HP Vertica Using the Command Line

Although HP supports installation on one node, two nodes, and multiple nodes, this section describes how to install the HP Vertica software on a cluster of nodes. It assumes that you have already performed the tasks in **Before You Install HP Vertica** (page [12](#)), and that you have an HP Vertica license key.

To install HP Vertica, follow the complete install procedure for both single-node and multi-node installations.

1 Download and install the HP Vertica install package (page [34](#))

2 Run the install script (page [35](#))

Special notes

- Downgrade installations are not supported.
- Be sure that you download the RPM for the correct operating system and architecture.
- HP Vertica supports two-node clusters with zero fault tolerance (K=0 safety). This means that you can add a node to a single-node cluster, as long as the installation node (the node upon which you build) is not the loopback node (`localhost/127.0.0.1`).
- You can also remove a node from a K-safe cluster.

Back Up Existing Databases

If you are doing an upgrade installation, back up the following for all existing databases:

- The Catalog and Data directories, using the HP Vertica backup utility. See Backing Up and Restoring the Database in the Administrator's Guide.
- `/opt/vertica/`, using manual methods. For example:

1. Enter the command:

```
tar -czvf /tmp/vertica.tgz /opt/vertica
```

2. Copy the tar file to a backup location.

Backing up MC

Before you **upgrade MC** (page [76](#)), HP recommends that you back up your MC metadata (configuration and user settings) on a storage location external to the server on which you installed MC.

- 1 On the target server (where you want to store MC metadata), log on as root or a user with sudo privileges.
- 2 Create a backup directory; for example:

```
# mkdir /backups/mc/mc-backup-20130425
```

- 3 Copy the `/opt/vconsole` directory to the new backup folder:

```
# cp -r /opt/vconsole /backups/mc/mc-backup-20130425
```

Download and Install the HP Vertica Install Package

To download and install the HP Vertica install package:

- 1 Use a Web browser to log in to **myVertica portal** <http://my.vertica.com/>.
- 2 Click the Download tab and download the HP Vertica install package to the Administration Host.

Be sure the package you download matches the operating system and the machine architecture on which you intend to install it. In the event of a node failure, you can use any other node to run the Administration Tools later.

- 3 If you installed a previous version of HP Vertica on any of the hosts in the cluster, use the Administration Tools to shut down any running database.

The database must stop normally; you cannot upgrade a database that requires recovery.

- 4 If you are using sudo, skip to the next step. If you are root, log in to the Administration Host as root (or log in as another user and switch to root).

```
$ su - root
password: root-password
#
```

Caution: When installing HP Vertica using an existing user as the dba, you must exit all UNIX terminal sessions for that user after setup completes and log in again to ensure that group privileges are applied correctly.

After HP Vertica is installed, you no longer need root privileges. See also **Verify sudo** (page 20).

- 5 Use one of the following commands to run the RPM package installer:

- If you are root and installing an RPM:
`rpm -Uvh pathname`
- If you are using sudo and installing an RPM:
\$ `sudo rpm -Uvh pathname`
- If you are using Debian, replace `rpm -Uvh` with `dpkg -i`

where *pathname* is the HP Vertica package file you downloaded.

Note: If the RPM installer reports multiple dependency problems, or you receive the error "ERROR: You're attempting to install the wrong RPM for this operating system", then you are trying to install the wrong HP Vertica install package. Make sure that the machine architecture (32-bit or 64-bit) of the package you downloaded matches the operating system.

After entering the command, a progress indicator appears:

```
Preparing...
##### [100%]
```

```

1:vertica
##### [100%]

HP Vertica 6.1.x.xx successfully installed on host hostname.
Additional install information may also appear.

```

Run the Install Script

The installation script takes the following as basic parameters:

- A list of hosts on which to install.
- The HP Vertica RPM path and filename.
- [Optionally] a username for the HP Vertica administrator. The install script creates a new system user-account. (dbadmin is used if you do not provide a username.)

Note: If you run the HP Vertica installation script without parameters, the script performs a single-node install on `localhost`. See *Installing HP Vertica Silently* (page 48) for details.

```
# /opt/vertica/sbin/install_vertica -s host_list -r rpm_package -u dba_username
```

<code>-s <i>host_list</i></code>	<p>A comma-separated list of hostnames or IP addresses to include in the cluster; do not include space characters in the list. Examples:</p> <pre>-s host01,host02,host03</pre> <pre>-s 192.168.233.101,192.168.233.102,192.168.233.103</pre> <p>This parameter is required on multi-node installations. On single-node installations the parameter is optional and the default is <i>localhost</i>. However, if you plan to expand to additional hosts later, then you must use the <code>-s</code> parameter. Installations done without the <code>-s</code> parameter cannot be upgraded to a multiple-node deployment, and must be reinstalled to support additional nodes.</p> <p>Note: If you are upgrading an existing installation of HP Vertica, specify one hostname, and HP Vertica finds the rest of the nodes.</p>
<code>-r <i>rpm_package</i></code>	<p>The pathname of the HP Vertica RPM package. Example:</p> <pre>-r "vertica_6.1.x.x.x86_64.RHEL5.rpm"</pre> <p>For Debian installs, provide the name of the Debian package, for example <code>-r vertica_5.x.xx-0_amd64.deb</code></p> <p>This parameter is required on multi-node installations only.</p>
<code>-u <i>dba_username</i></code>	<p>The name of the Database Administrator Linux account to create. Only this account can run the Administration Tools. If you omit the <code>-u</code> parameter, the default database administrator account name is <code>dbadmin</code>.</p> <p>This parameter is optional for new installations done as root but must be specified when upgrading or when installing using <code>sudo</code>. If upgrading, use the <code>-u</code> parameter to specify the same DBA account name that you used previously. If installing using <code>sudo</code>, the user must already exist.</p> <p>Note: If you manually create the user, modify the user's <code>.bashrc</code> file to</p>

	include the line: <code>PATH=/opt/vertica/bin:\$PATH</code>
--	--

The complete set of **installation parameters** (page [37](#)), including optional parameters and example output, is described in `install_vertica` (page [37](#)).

When run on a cluster (more than one node), the install script installs the provided RPM on all nodes and sets up passwordless ssh for the administrator user across all the hosts. (If passwordless ssh was already setup, it verifies that it is functioning correctly. If it is already set up but not functioning correctly, it backs up the existing keys to `.ssh.vtbackup` and sets up new keys.) It checks network connectivity across the hosts.

On a single-node localhost installation, the installer does not set up passwordless ssh. Because HP Vertica's backup scripts require that the administrator can log into the node via ssh without a password, you need to manually enable passwordless ssh logins for any single-node install if you want to use the backup scripts.

The install script also creates an operating system user called `spread`, which is responsible for running the spread daemon. User `spread` is non-configurable and is set to `no login`.

The install script modifies some operating system configuration settings and warns about issues with the setup that could prevent HP Vertica from functioning properly.

To perform a typical install of HP Vertica:

1 Run the install script:

Note: You must run the install script using the BASH shell. Other shells are not supported.

- If you are root:

```
# /opt/vertica/sbin/install_vertica -s host_list -r rpm_package -u dba_username
```

- If you are using sudo:

```
$ sudo /opt/vertica/sbin/install_vertica -s host_list -r rpm_package -u dba_username
```

2 When prompted, provide the requested password to allow installation of the RPM and system configuration of the other cluster nodes. If you are root, this is the root password. If you are using sudo, this is the sudo user password. The password does not echo on the command line.

For example, root would see:

```
HP Vertica Database 6.1.x Installation Tool
Please enter password for root@host01:password
```

3 Provide the database administrator account password when prompted:

```
Enter password for new UNIX user dbadmin:password
Retype new UNIX password for user dbadmin:password
```

4 Carefully examine any warnings produced by `install_vertica` and correct the problems, if possible. For example, insufficient RAM, insufficient network throughput, and too high readahead settings on the filesystem could cause performance problems later on. Additionally, LANG warnings, if not resolved, can cause database startup to fail and issues with VSQL. The system LANG attributes must be UTF-8 compatible.

5 Disconnect from the Administration Host as instructed by the script. This is required to:

- Set certain system parameters correctly.
- Function as the HP Vertica database administrator.

At this point, Linux root privileges are no longer needed. The database administrator can perform the remaining steps.

Exception: when creating a new database, the database administrator might want to use different data or catalog locations than those created by the installation script. In that case, a Linux administrator might need to create those directories and change their ownership to the database administrator.

- 6 Log in to the Database Administrator account on the administration host.
- 7 If you are installing the Enterprise Edition product, you must install the license key you downloaded previously as described in ***Install the License Key*** (page [63](#)).
- 8 **Optionally** perform the following steps:
 - Install the ODBC and JDBC driver on non-cluster hosts. The ODBC and JDBC drivers are installed by default on cluster hosts, but still must be configured. See the Programmer's Guide for details.
 - ***Install the vsq client application on non-cluster hosts*** (page [65](#)).
- 9 If you have not already done so, proceed to the Tutorial in the Getting Started Guide. Otherwise, proceed to Configuring the Database in the Administrator's Guide.

About the install_vertica Script

The `install_vertica` script creates a cluster of nodes on which you can create a database. The script does a number of checks to catch common Linux misconfigurations, such as checking the connectivity and bandwidth characteristics of the communication links among the cluster machines and verifying other prerequisites.

Note: The `install_vertica` requires the calling user's shell to be `/bin/bash`. Other shells give unpredictable results and are not supported.

Syntax

```
/opt/vertica/sbin/install_vertica
-s host_list
-r package_name
[ -d data_directory ]
[ -g user_group ]
[ -l dba_home_directory ]
[ -p dba_password ]
[ -w ]
[ -P root_password ]
[ -i file ]
[ -u dba_username ]
[ -z file ]
[ -A hosts_to_add ]
[ -B file ]
[ -C ]
[ -E ]
[ -L { license_file | CE } ]
```

```
[ -N ]
[ -Q ]
[ -R ]
[ -S ]
[ -T ]
[ -U ]
[ -Y ]
```

Parameters

If you run `install_vertica` without parameters, the script performs a single-node install on localhost. See **Installing HP Vertica silently** (page [48](#)) for details.

<code>-s host_list</code>	<p>A comma-separated list of hostnames or IP addresses to include in the cluster. Do not include space characters in the list. Examples:</p> <pre>-s host01,host02,host03 -s 192.168.233.101,192.168.233.102,192.168.233.103</pre> <p>The <code>-s</code> parameter is required on multi-node installations only. On single-node installations the default is <code>localhost</code>.</p> <p>Note: If you are upgrading an existing installation of HP Vertica, be sure to use the same host names that you used previously.</p>
<code>-r package_name</code>	<p>The name of the RPM or Debian package that contained this script. Example:</p> <pre>-r vertica_6.1.x.x.x86_64.RHEL5.rpm</pre> <p>The <code>-r</code> parameter is required on multi-node installations only. It has no default value.</p>
<code>-d data_directory</code>	<p>The default directory for database data and catalog files. The default is <code>/home/dbadmin</code></p> <p>Note: Do not use a shared directory over more than one host for this setting. Data and catalog directories must be distinct for each node. Multiple nodes must not be allowed to write to the same data or catalog directory.</p>
<code>-g user_group</code>	<p>The unix group for DBA users. The default is <code>verticadba</code>.</p>
<code>-l dba_home_directory</code>	<p>The home directory for the database administrator. The default is <code>/home/dbadmin</code></p>
<code>-p dba_password</code>	<p>The password for the database administrator account. If not supplied, the script prompts for a password and does not echo the input.</p>
<code>-w</code>	<p>Configures spread to output logging output to <code>/opt/vertica/log/spread_<hostname>.log</code> (default).</p>
<code>-P root_password</code>	<p>The root password to use by default for each cluster host. If not supplied, and the <code>-i</code> option is not used, then the script prompts for the password if and when necessary and does not echo the input. Do not use with the <code>-i</code> option.</p> <p>Special note about root password:</p> <p>If you run the <code>install_vertica</code> script as root, specify the root password with the <code>-P</code> parameter:</p> <pre># /opt/vertica/sbin/install_vertica -P <root_passwd></pre> <p>If, however, you run the <code>install_vertica</code> script with the <code>sudo</code> command, the</p>

	<p>password for the <code>-P</code> parameter should be the password of the user who runs <code>install_vertica</code>, not the root password. For example if user <code>dbadmin</code> runs <code>install_vertica</code> with <code>sudo</code> and has a password with the value <code>dbapasswd</code>, then the value for <code>-P</code> should be <code>dbapasswd</code>:</p> <pre>\$ sudo /opt/vertica/sbin/install_vertica -P dbapasswd</pre>
<code>-i file</code>	<p>The root private-key <i>file</i> to use if passwordless ssh has already been configured between the hosts. Verify that normal SSH works without a password before using this option. The file can be private key file (for example, <code>id_rsa</code>), or PEM file. Do not use with the <code>-P</code> option.</p> <p>Note: You cannot run <code>install_vertica</code> with the <code>sudo</code> command when using this parameter.</p>
<code>-u dba_username</code>	<p>The name of the database administrator account to create. Only this account can run the Administration Tools. If you omit the <code>-u</code> parameter, the default database administrator account name is <code>dbadmin</code>.</p> <p>Note: This parameter is optional for new installations done as root but must be specified when upgrading or when installing using <code>sudo</code>. If upgrading, use the <code>-u</code> parameter to specify the same DBA account name that you used previously. If installing using <code>sudo</code>, the user must already exist. Otherwise you might encounter problems later.</p>
<code>-z file</code>	<p>Accepts an existing properties file created by <code>-B file_name</code>. This properties file contains key/value parameters that map to values in the <code>install_vertica</code> script, many with boolean arguments that default to false.</p>
<code>-A</code>	<p>A comma-separated list of hosts to add to an existing HP Vertica cluster.</p> <p><code>-A</code> modifies an existing installation of HP Vertica by adding a host to the database cluster and then reconfiguring the spread. This is useful for increasing system performance or setting K-safety to one (1) or two (2).</p> <p>Notes:</p> <ul style="list-style-type: none"> ▪ If you have used the <code>-T</code> parameter to configure spread to use direct point-point communication within the existing cluster, you must use the <code>-T</code> parameter when you add a new host; otherwise, the new host automatically uses UDP broadcast traffic, resulting in cluster communication problems that prevent HP Vertica from running properly. Examples: ▪ <code>-A host01</code> ▪ <code>-A 192.168.233.101</code> ▪ The <code>update_vertica</code> script described in Adding Nodes calls the <code>install_vertica</code> script to update the installation. You can use either the <code>install_vertica</code> or <code>update_vertica</code> script with the <code>-A</code> parameter.
<code>-B file_name</code>	<p>Accepts a file name, which when used in conjunction with command line options, creates a properties file that can be used with the <code>-z</code> parameter. This parameter creates the properties file and exits; it has no impact on installation.</p>
<code>-C</code>	<p>Cleans previously stored configuration files if the database is not defined. Use this parameter if you need to change the hosts that are included in your cluster. You can use this parameter only when no database is already defined.</p>

-E	<p>Allows a new node to replace a failed node in the database. Without this parameter, HP Vertica does not allow the failed node to be dropped because it is still considered in use by the database even though it has failed.</p> <p>This parameter must be used in combination with the -A (add) and -R (remove) host parameters. For example:</p> <pre>-A 192.168.233.102 -R 192.168.233.101 -E</pre> <p>Note: After running <code>install_vertica</code> with -A -R -E to replace a failed host, run <code>install_vertica</code> again with just -R after you have removed the node from the database. This action clears the node that you replaced from the Administration Tools metadata.</p>
-L { <i>license_file</i> CE }	<p>Silently and automatically deploys the license key to <code>/opt/vertica/config/share</code>. On multi-node installations, the -L option also applies the license to all nodes declared in the -s <code>host_list</code>.</p> <p>If specified with CE, automatically deploys the Community Edition license key, which is included in your download. You do not need to specify a license file.</p>
-N	<p>Indicates that HP Vertica is installed on hosts that exist on different subnets.</p>
-Q	<p>Ends setup without asking the user to run the Administration Tools.</p>
-R	<p>A comma-separated list of hosts to remove from an existing HP Vertica cluster.</p> <p>-R modifies an existing installation of HP Vertica by removing a host from the database cluster and then reconfiguring the spread. This is useful for removing an obsolete or over-provisioned system. For example:</p> <pre>-R host01 -R 192.168.233.101</pre> <p>Notes:</p> <ul style="list-style-type: none"> If you have identified a node using various forms of the host name and IP address, you must identify all the forms you used. For example, you might identify a node with both short and fully-qualified names. Use a comma-separated list to identify two or more forms of the host name. Do not add a space between the commas. For example: <pre>-R host01,Host01.vertica.com</pre> If you used the -T parameter to configure spread to use direct point-point communication within the existing cluster, you must use -T when you remove a host; otherwise, the hosts automatically use UDP broadcast traffic, resulting in cluster communication problems that prevents HP Vertica from running properly. The <code>update_vertica</code> script described in Removing Nodes in the Administrator's Guide calls the <code>install_vertica</code> script to perform the update to the installation. You can use either the <code>install_vertica</code> or <code>update_vertica</code> script with the -R parameter.
-S	<p>Takes either the value 'default' or a broadcast network IP address to allow spread communications to be configured on a subnet that is different from other HP Vertica data communications. -S is also used to force a cluster-wide spread reconfiguration when changing spread related options—it is needed when you switch between -T and U. -S is required when adding or replacing a host.</p> <p>Important: Stop all databases before you use the -S parameter. If the databases are not stopped then the -S parameter causes the database to restart.</p>

-T	Configures spread to use direct point-to-point communication between all HP Vertica nodes. You should use this option if your nodes aren't located on the same subnet. You should also use this option for all virtual environment installations, regardless of whether the virtual servers are on the same subnet or not. Important: When changing the configuration from -U (the default) to -T or from -T to -U, the -S parameter must also be used.
-U	Specifies that HP Vertica use UDP broadcast traffic by spread between nodes on the subnet. This parameter is automatically used by default. Important: When changing the configuration from -U (the default) to -T or from -T to -U, the -S parameter must also be used.
-Y	Silently accepts the EULA agreement. On multi-node installations, the -Y value is propagated throughout the cluster at the end of the installation, at the same time as the Administration Tools metadata.

Notes

Do not be concerned by the following message:

```
Checking ...Failed.
```

It indicates that `install_vertica` found and is correcting an incorrect configuration parameter.

Example

The following example creates a cluster of nodes (host01, host02, and host03) using the `vertica_6.1.x.x.x86_64.SUSE10 RPM` package.

Note: The output is for illustrative purposes only and may change slightly from release to release.

```
# /opt/vertica/sbin/install_vertica -s fc10-1,fc10-2,fc10-3,fc10-4 -r
vertica_6.1.x.x86_64.SUSE10.rpm
Vertica Analytic Database 6.1.x.x Installation Tool
Enter password for root@host01 (2 attempts left):
root@host01: -- Checking status of ntpd service -- ...Done
root@host02: -- Checking status of ntpd service -- ...Done
root@host03: -- Checking status of ntpd service -- ...Done
root@host04: -- Checking status of ntpd service -- ...Done
root@host01: -- Checking ntp stratum level -- ...Done
root@host02: -- Checking ntp stratum level -- ...Done
root@host03: -- Checking ntp stratum level -- ...Done
root@host04: -- Checking ntp stratum level -- ...Done
root@host01: -- Checking readahead parameter for I/O devices -- ...Done
root@host02: -- Checking readahead parameter for I/O devices -- ...Done
root@host03: -- Checking readahead parameter for I/O devices -- ...Done
root@host04: -- Checking readahead parameter for I/O devices -- ...Done
root@host01: -- Checking vm.min_free_kbytes parameter in /sbin/sysctl -- ...Done
root@host02: -- Checking vm.min_free_kbytes parameter in /sbin/sysctl -- ...Done
root@host03: -- Checking vm.min_free_kbytes parameter in /sbin/sysctl -- ...Done
root@host04: -- Checking vm.min_free_kbytes parameter in /sbin/sysctl -- ...Done
root@host01: -- Checking maximum number of open file descriptors in /etc/security/limits.conf --
...Done
root@host02: -- Checking maximum number of open file descriptors in /etc/security/limits.conf --
...Done
root@host03: -- Checking maximum number of open file descriptors in /etc/security/limits.conf --
...Done
root@host04: -- Checking maximum number of open file descriptors in /etc/security/limits.conf --
```

Installation Guide

```
...Done
root@host01: -- Executing -- sudo /usr/sbin/groupadd verticadba
root@host02: -- Executing -- sudo /usr/sbin/groupadd verticadba
root@host03: -- Executing -- sudo /usr/sbin/groupadd verticadba
root@host04: -- Executing -- sudo /usr/sbin/groupadd verticadba
root@host01: -- Executing -- sudo /usr/bin/id -nu dbadmin
root@host02: -- Executing -- sudo /usr/bin/id -nu dbadmin
root@host03: -- Executing -- sudo /usr/bin/id -nu dbadmin
root@host04: -- Executing -- sudo /usr/bin/id -nu dbadmin
root@host01: -- Executing -- sudo sh -c "chown dbadmin:dbadmin ~dbadmin"
root@host04: -- Executing -- sudo sh -c "chown dbadmin:dbadmin ~dbadmin"
root@host02: -- Executing -- sudo sh -c "chown dbadmin:dbadmin ~dbadmin"
root@host03: -- Executing -- sudo sh -c "chown dbadmin:dbadmin ~dbadmin"
root@host01: -- Executing -- sudo sh -c "chmod 700 ~dbadmin"
root@host04: -- Executing -- sudo sh -c "chmod 700 ~dbadmin"
root@host02: -- Executing -- sudo sh -c "chmod 700 ~dbadmin"
root@host03: -- Executing -- sudo sh -c "chmod 700 ~dbadmin"
root@host01: -- Adding user dbadmin to dbagroup -- ...Done
root@host02: -- Adding user dbadmin to dbagroup -- ...Done
root@host03: -- Adding user dbadmin to dbagroup -- ...Done
root@host04: -- Adding user dbadmin to dbagroup -- ...Done
root@host01: -- Checking /scratch_b/qa directory -- ...Done
root@host02: -- Checking /scratch_b/qa directory -- ...Done
root@host03: -- Checking /scratch_b/qa directory -- ...Done
root@host04: -- Checking /scratch_b/qa directory -- ...Done
root@host01: -- Testing permissions -- ...Done
root@host02: -- Testing permissions -- ...Done
root@host03: -- Testing permissions -- ...Done
root@host04: -- Testing permissions -- ...Done
root@host01: -- Changing the owner of /scratch_b/qa directory to dbadmin -- ...Done
root@host02: -- Changing the owner of /scratch_b/qa directory to dbadmin -- ...Done
root@host03: -- Changing the owner of /scratch_b/qa directory to dbadmin -- ...Done
root@host04: -- Changing the owner of /scratch_b/qa directory to dbadmin -- ...Done
root@host01: -- Write check for /scratch_b/qa directory -- ...Done
root@host02: -- Write check for /scratch_b/qa directory -- ...Done
root@host03: -- Write check for /scratch_b/qa directory -- ...Done
root@host04: -- Write check for /scratch_b/qa directory -- ...Done ===== Host Availability Test
=====
All hosts are available ... ok

===== Testing fc10-2.vertica.com =====
Host availability ... ok
Passwordless SSH access ... ok
Passwordless SSH access to other hosts ... ok
Getting various information about the host ... ok
Python is installed ... ok
Enough RAM per CPUs ... ok
Getting information about the ports ... ok
Port availability for Spread ... ok
Port availability for Vertica ... ok
Gathering information about the network ... ok
Sanity of /etc/hosts ... ok
Hostname resolution ... ok
IP configuration ... ok

===== Testing fc10-1.vertica.com =====
Host availability ... ok
Passwordless SSH access ... ok
Passwordless SSH access to other hosts ... ok
Getting various information about the host ... ok
Python is installed ... ok
Enough RAM per CPUs ... ok
Getting information about the ports ... ok
Port availability for Spread ... ok
Port availability for Vertica ... ok
Gathering information about the network ... ok
Sanity of /etc/hosts ... ok
Hostname resolution ... ok
```

```

IP configuration                               ... ok

==== Testing fc10-3.vertica.com =====
Host availability                             ... ok
Passwordless SSH access                      ... ok
Passwordless SSH access to other hosts       ... ok
Getting various information about the host     ... ok
Python is installed                          ... ok
Enough RAM per CPUs                          ... ok
Getting information about the ports            ... ok
Port availability for Spread                  ... ok
Port availability for Vertica                  ... ok
Gathering information about the network        ... ok
Sanity of /etc/hosts                          ... ok
Hostname resolution                           ... ok
IP configuration                               ... ok

==== Testing fc10-4.vertica.com =====
Host availability                             ... ok
Passwordless SSH access                      ... ok
Passwordless SSH access to other hosts       ... ok
Getting various information about the host     ... ok
Python is installed                          ... ok
Enough RAM per CPUs                          ... ok
Getting information about the ports            ... ok
Port availability for Spread                  ... ok
Port availability for Vertica                  ... ok
Gathering information about the network        ... ok
Sanity of /etc/hosts                          ... ok
Hostname resolution                           ... ok
IP configuration                               ... ok

==== Consistency Test =====
IP and Netmask consistency                    ... ok
All hosts are on the same subnet               ... ok
LANG and TZ environment variables              ... ok

==== Network Test =====
Network communication                          ... ok

Test of host fc10-2.vertica.com (ok) =====
    Hostname resolution (ok)
    -----
    Info: The $HOSTNAME environment variable (fc10-2) does not match the actual hostname
    (fc10-2.vertica.com)

Test of host fc10-1.vertica.com (ok) =====
    Hostname resolution (ok)
    -----
    Info: The $HOSTNAME environment variable (fc10-1) does not match the actual hostname
    (fc10-1.vertica.com)

Test of host fc10-3.vertica.com (ok) =====
    Hostname resolution (ok)
    -----
    Info: The $HOSTNAME environment variable (fc10-3) does not match the actual hostname
    (fc10-3.vertica.com)

Test of host fc10-4.vertica.com (ok) =====
    Hostname resolution (ok)

```

Installation Guide

```
-----
Info: The $HOSTNAME environment variable (fc10-4) does not match the actual hostname
(fc10-4.vertica.com)
```

```
Network Test (ok)
=====
```

```
-----
Network communication (ok)
-----
Low throughput fc10-1.vertica.com to fc10-2.vertica.com: 84.2610416353 Mbps; check network
interface/switch configuration
Low throughput fc10-3.vertica.com to fc10-2.vertica.com: 87.9604060062 Mbps; check network
interface/switch configuration
Low throughput fc10-4.vertica.com to fc10-2.vertica.com: 87.0131137763 Mbps; check network
interface/switch configuration
Low throughput fc10-2.vertica.com to fc10-1.vertica.com: 87.9564862945 Mbps; check network
interface/switch configuration
Low throughput fc10-3.vertica.com to fc10-1.vertica.com: 86.9869472003 Mbps; check network
interface/switch configuration
Low throughput fc10-4.vertica.com to fc10-1.vertica.com: 85.1561844917 Mbps; check network
interface/switch configuration
Low throughput fc10-2.vertica.com to fc10-3.vertica.com: 87.5655859204 Mbps; check network
interface/switch configuration
Low throughput fc10-1.vertica.com to fc10-3.vertica.com: 87.5760989283 Mbps; check network
interface/switch configuration
Low throughput fc10-4.vertica.com to fc10-3.vertica.com: 87.5283536148 Mbps; check network
interface/switch configuration
Low throughput fc10-2.vertica.com to fc10-4.vertica.com: 85.1868837143 Mbps; check network
interface/switch configuration
Low throughput fc10-1.vertica.com to fc10-4.vertica.com: 87.0131137763 Mbps; check network
interface/switch configuration
Low throughput fc10-3.vertica.com to fc10-4.vertica.com: 87.9585613925 Mbps; check network
interface/switch configuration
```

```
root@host01: -- Stopping spread daemon -- ...Done
root@host03: -- Stopping spread daemon -- ...Done
root@host04: -- Stopping spread daemon -- ...Done
root@host01: -- Removing spread daemon -- ...Done
root@host03: -- Removing spread daemon -- ...Done
root@host04: -- Removing spread daemon -- ...Done
root@host01: -- Removing old spread daemon links -- ...Done
root@host04: -- Removing old spread daemon links -- ...Done
root@host02: -- Removing old spread daemon links -- ...Done
root@host03: -- Removing old spread daemon links -- ...Done
root@host01: -- Setting up spread daemon -- ...Done
root@host04: -- Setting up spread daemon -- ...Done
root@host02: -- Setting up spread daemon -- ...Done
root@host03: -- Setting up spread daemon -- ...Done
root@host01: -- Creating spread configuration file -- ...Done
root@host04: -- Creating spread configuration file -- ...Done
root@host02: -- Creating spread configuration file -- ...Done
root@host03: -- Creating spread configuration file -- ...Done
root@host01: -- Executing -- sudo cat /opt/vertica/config/vspread.conf
root@host04: -- Executing -- sudo cat /opt/vertica/config/vspread.conf
root@host02: -- Executing -- sudo cat /opt/vertica/config/vspread.conf
root@host03: -- Executing -- sudo cat /opt/vertica/config/vspread.conf
root@host01: -- Checking /opt/vertica/config directory -- ...Done
root@host04: -- Checking /opt/vertica/config directory -- ...Done
root@host02: -- Checking /opt/vertica/config directory -- ...Done
root@host03: -- Checking /opt/vertica/config directory -- ...Done
root@host01: -- Testing permissions -- ...Done
root@host04: -- Testing permissions -- ...Done
root@host02: -- Testing permissions -- ...Done
root@host03: -- Testing permissions -- ...Done
root@host01: -- Write check for /opt/vertica/config directory -- ...Done
root@host04: -- Write check for /opt/vertica/config directory -- ...Done
```


Installation Guide

[illegible]

```

root@host03: -- Checking /opt/vertica/config/users directory -- ...Done
root@host04: -- Checking /opt/vertica/config/users directory -- ...Done
root@host01: -- Testing permissions -- ...Done
root@host02: -- Testing permissions -- ...Done
root@host03: -- Testing permissions -- ...Done
root@host04: -- Testing permissions -- ...Done
root@host01: -- Write check for /opt/vertica/config/users directory -- ...Done
root@host02: -- Write check for /opt/vertica/config/users directory -- ...Done
root@host03: -- Write check for /opt/vertica/config/users directory -- ...Done
root@host04: -- Write check for /opt/vertica/config/users directory -- ...Done
root@host01: -- Changing the group of /opt/vertica/config/users directory to verticadba -- ...Done
root@host02: -- Changing the group of /opt/vertica/config/users directory to verticadba -- ...Done
root@host03: -- Changing the group of /opt/vertica/config/users directory to verticadba -- ...Done
root@host04: -- Changing the group of /opt/vertica/config/users directory to verticadba -- ...Done
root@host01: -- Checking /opt/vertica/log directory -- ...Done
root@host02: -- Checking /opt/vertica/log directory -- ...Done
root@host03: -- Checking /opt/vertica/log directory -- ...Done
root@host04: -- Checking /opt/vertica/log directory -- ...Done
root@host01: -- Testing permissions -- ...Done
root@host02: -- Testing permissions -- ...Done
root@host03: -- Testing permissions -- ...Done
root@host04: -- Testing permissions -- ...Done
root@host01: -- Write check for /opt/vertica/log directory -- ...Done
root@host02: -- Write check for /opt/vertica/log directory -- ...Done
root@host03: -- Write check for /opt/vertica/log directory -- ...Done
root@host04: -- Write check for /opt/vertica/log directory -- ...Done
root@host01: -- Changing the group of /opt/vertica/log directory to verticadba -- ...Done
root@host02: -- Changing the group of /opt/vertica/log directory to verticadba -- ...Done
root@host03: -- Changing the group of /opt/vertica/log directory to verticadba -- ...Done
root@host04: -- Changing the group of /opt/vertica/log directory to verticadba -- ...Done
root@host01: -- Testing permissions -- ...Done
root@host02: -- Testing permissions -- ...Done
root@host03: -- Testing permissions -- ...Done
root@host04: -- Testing permissions -- ...Done
root@host01: -- Checking /opt/vertica/config/users/dbadmin directory -- ...Done
root@host02: -- Checking /opt/vertica/config/users/dbadmin directory -- ...Done
root@host03: -- Checking /opt/vertica/config/users/dbadmin directory -- ...Done
root@host04: -- Checking /opt/vertica/config/users/dbadmin directory -- ...Done
root@host01: -- Testing permissions -- ...Done
root@host02: -- Testing permissions -- ...Done
root@host03: -- Testing permissions -- ...Done
root@host04: -- Testing permissions -- ...Done
root@host01: -- Changing the owner of /opt/vertica/config/users/dbadmin directory to dbadmin --
...Done
root@host02: -- Changing the owner of /opt/vertica/config/users/dbadmin directory to dbadmin --
...Done
root@host03: -- Changing the owner of /opt/vertica/config/users/dbadmin directory to dbadmin --
...Done
root@host04: -- Changing the owner of /opt/vertica/config/users/dbadmin directory to dbadmin --
...Done
root@host01: -- Write check for /opt/vertica/config/users/dbadmin directory -- ...Done
root@host02: -- Write check for /opt/vertica/config/users/dbadmin directory -- ...Done
root@host03: -- Write check for /opt/vertica/config/users/dbadmin directory -- ...Done
root@host04: -- Write check for /opt/vertica/config/users/dbadmin directory -- ...Done Error Monitor
0 errors 0 warnings Installation complete.

```

To create a database:

1. Logout and login as dbadmin.
2. Run /opt/vertica/bin/adminTools as dbadmin
3. Select Create Database from the Configuration Menu

To ensure that group privileges are properly applied, you will need to log out of any existing sessions for the user dbadmin and login again.

Installing HP Vertica Silently

This section describes how to create a properties file that lets you install and deploy HP Vertica-based applications quickly and without much manual intervention.

Note: The procedure assumes that you have already performed the tasks in *Before you Install HP Vertica* (page [12](#)).

Installing HP Vertica silently

- 1 Download and install the HP Vertica install package, as described in *Installing HP Vertica* (page [33](#)).
- 2 Create the properties file that enables non-interactive setup by supplying the parameters you want HP Vertica to use. For example:

The following command assumes a multi-node setup:

```
# /opt/vertica/sbin/install_vertica -B file_name -L /tmp/license.txt -Y
-p password -P password -s host_list -r package_name
```

The following command assumes a single-node setup:

```
# /opt/vertica/sbin/install_vertica -B file_name -L /tmp/license.txt -Y
-p password -P password
```

For single-node installations, the `-s` and `-r` parameters are not required.

<code>-B file_name</code>	[Required] Accepts a file name, which when used in conjunction with command line options, creates a properties file that can be used with the <code>-z</code> option during setup. This flag creates the properties file and exits; it has no impact on installation.
<code>-L { license_file CE }</code>	Silently and automatically deploys the license key to <code>/opt/vertica/config/share</code> . On multi-node installations, the <code>-L</code> option also applies the license to all nodes declared in the <code>-s host_list</code> . If specified with <code>CE</code> , automatically deploys the Community Edition license key, which is included in your download. You do not need to specify a license file.
<code>-Y</code>	Silently accepts the EULA agreement during setup.
<code>-p dba_password</code>	The password for the Database Administrator account; if not supplied, the script prompts for the password and does not echo the input.
<code>-P root_password</code>	The root password to use by default for each cluster host; if not supplied, the script prompts for the password if and when necessary and does not echo the input.
<code>-s host_list</code>	A comma-separated list of hostnames or IP addresses to include in the cluster; do not include space characters in the list. Examples: <pre>-s host01,host02,host03</pre> <pre>-s 192.168.233.101,192.168.233.102,192.168.233.103</pre> <p>This parameter is required on multi-node installations only. On single-node installations the default is <code>localhost</code>. If you are upgrading an existing installation of HP Vertica, be sure to</p>

	use the same host names.
<code>-r package_name</code>	<p>The name of the RPM or Debian package that contained this script. Example:</p> <pre>-r vertica_6.1.x.x.x86_64.RHEL5.rpm</pre> <p>This parameter is required on multi-node installations only. It has no default value.</p>

See **The *install_vertica* Script** (page [37](#)) for the complete set of installation parameters.

TIP:

Supply the parameters to the properties file once only. You can then install HP Vertica using just the `-z` parameters, as described in the Step 3.

- Use one of the following commands to run the installation script.
 - If you are root:


```
# /opt/vertica/sbin/install_vertica -z file_name
```
 - If you are using sudo:


```
$ sudo /opt/vertica/sbin/install_vertica -z file_name
```

`-z file_name` Accepts an existing properties file created by `-B file_name`. This properties file contains key/value parameters that map to values in the `install_vertica` script, many with boolean arguments that default to false

The command for a single-node install might look like this:

```
# /opt/vertica/sbin/install_vertica -z /tmp/vertica-inst.prp
```
- If you did not supply a `-P` password parameter to the properties file, you are prompted to provide the requested password to allow installation of the RPM and system configuration of the other cluster nodes. If you are root, this is the root password. If you are using sudo, this is the sudo user password. The password does not echo on the command line.

Note: If you are root on a single-node installation, you are not prompted for a password.
- If you did not supply a `-p` password parameter to the properties file, you are prompted to provide the database administrator account password.

The installation script creates a new Linux user account (dbadmin by default) with the password that you provide.
- Carefully examine any warnings produced by `install_vertica` and correct the problems if possible. For example, insufficient RAM, insufficient Network throughput and too high readahead settings on filesystem could cause performance problems later on.

Ignore insufficient RAM warnings when installing on VMware.

Note: You can redirect any warning outputs to a separate file, instead of having them display on the system. Use the `redirect_output = filename` parameter.
- Optionally perform the following steps:
 - Install the ODBC and JDBC driver.
 - Install the vsqI client application on non-cluster hosts** (page [65](#)).
- Disconnect from the Administration Host as instructed by the script. This is required to:
 - Set certain system parameters correctly.

- Function as the HP Vertica database administrator.

At this point, Linux root privileges are no longer needed. The database administrator can perform the remaining steps.

Exception: when creating a new database, the database administrator might want to use different data or catalog locations than those created by the installation script. In that case, a Linux administrator might need to create those directories and change their ownership to the database administrator.

- 7 If you supplied the `-L` parameter to the properties file, skip to step 11. Otherwise, log in to the Database Administrator account on the administration host.
- 8 Install the license key you downloaded previously as described in ***Install the license key*** (page [63](#)).
- 9 Proceed to the Tutorial in the Getting Started Guide and then see Configuring the database in the Administrator's Guide.

Notes

- Downgrade installations are not supported.
- If you installed HP Vertica on localhost, add a line in `/etc/hosts` that reads:
127.0.0.1 localhost
- The following are the values within the option parser and are the properties supported by `-z`:

<code>hosts</code>	<code>// comma separated lists of hosts</code>
<code>add_hosts</code>	<code>// comma separated list of hosts to add to</code>
<code>existing_cluster</code>	
<code>remove_hosts</code>	<code>// comma separated list of hosts to remove</code>
<code>from_cluster</code>	
<code>rpm_file_name</code>	
<code>vertica_dba_user</code>	<code>(default dbadmin)</code>
<code>vertica_dba_group</code>	<code>(default verticadba)</code>
<code>vertica_dba_user_password</code>	<code>// clear text!</code>
<code>root_password</code>	<code>// clear text!</code>
<code>dba_user_dir</code>	
<code>data_dir</code>	
<code>debug</code>	<code>// boolean</code>
<code>clean</code>	<code>// boolean</code>
<code>ignore_netmask</code>	<code>// boolean</code>
<code>allowUDP</code>	<code>// boolean</code>
<code>direct_only</code>	<code>// boolean</code>
<code>forceSpreadReconfiguration</code>	<code>// boolean</code>
<code>replaceHost</code>	<code>// boolean</code>
<code>skip_network_test</code>	<code>// boolean</code>
<code>license_file</code>	
<code>accept_eula</code>	<code>//boolean</code>
<code>redirect_output</code>	<code>// redirect stdout and stderr here!</code>

The following is an example of the contents of the properties file:

```
accept_eula = True
license_file = /tmp/license.txt
record_to = file_name
```

```
root_password = password
vertica_dba_group = verticadba
vertica_dba_user = dbadmin
vertica_dba_user_password = password
```

Installing HP Vertica on Amazon Web Services (AWS)

Beginning with Vertica 6.1.x, you can use Vertica on AWS by utilizing a pre-configured Amazon Machine Image (AMI). For details on installing and configuring a cluster on AWS, refer to the document, *Using HP Vertica on Amazon Web Services*.

Creating an HP Vertica Cluster Using MC

You can use Management Console to install an HP Vertica cluster on hosts where HP Vertica software has not been installed. The Cluster Installation wizard lets you specify the hosts you want to include in your HP Vertica cluster, loads the HP Vertica software onto the hosts, validates the hosts, and assembles the nodes into a cluster.

Management Console must be installed and configured before you can create a cluster on targeted hosts. See *Installing and Configuring the MC* (page [58](#)) for details.

Steps required to install an HP Vertica cluster using MC:

- **Install and configure MC** (page [58](#))
- **Prepare the hosts** (page [51](#))
- **Create the private key file** and copy it to your local machine
- **Run the Cluster Installation Wizard** (page [54](#))
- **Validate the hosts and create the cluster**
- **Create a new database on the cluster**

Prepare the Hosts

Before you can install an HP Vertica cluster using the MC, you must prepare each host that will become a node in the cluster. The cluster creation process runs validation tests against each host before it attempts to install the HP Vertica software. These tests ensure that the host is correctly configured to run HP Vertica.

The validation tests provide:

- Warnings and error messages when they detect a configuration setting that conflicts with HP Vertica's requirements or any performance issue
- Suggestions for configuration changes when they detect an issue

Note: The validation tests do not automatically fix problems they encounter.

All hosts must pass validation before the cluster can be created.

If you accepted the default configuration options when installing the OS on your host, then the validation tests will likely return errors, since some of the default options used on Linux systems conflict with HP Vertica requirements. See the Installation Guide for details on OS settings. To speed up the validation process you can perform the following steps on the prospective hosts before you attempt to validate the hosts. These steps are based on Red Hat Enterprise Linux and CentOS systems, but other supported platforms have similar settings.

How to prepare hosts

On each host you want to include in the HP Vertica cluster, you must:

- 1 Stage the host according to **Configuring Hardware and Installing Linux** (page [14](#)).
Also note that all the hosts should be running the same distribution and version of a supported Linux platform. See the **Supported Platforms Guide** (<http://myvertica.vertica.com>) for details.
- 2 Disable selinux:
 1. Edit `/etc/sysconfig/selinux` and change setting for SELINUX to disabled. This disables selinux at boot time.
 2. Type `setenforce 0` to disable selinux immediately.
- 3 Enable NTP:
 1. `/sbin/chkconfig ntpd on`
 2. `/sbin/service ntpd restart`
- 4 Set the timezone variable, if you do not have it set system-wide, by adding the following to the root account's `.bashrc` (Eastern US Time example):
`export TZ="America/New York"`
- 5 Modify `pam_limits.so` to allow the conveying of limits with the `su` command by editing `/etc/pam.d/su` and adding the following:
`session required pam_limits.so`
- 6 Disable CPU Frequency Scaling:
 1. `/sbin/chkconfig cpuspeed off`
 2. `/sbin/service cpuspeed stop`
- 7 Turn off hugepages defragmentation by adding the below lines to `/etc/rc.local` so they are active at boot time:
 1. `echo never > /sys/kernel/mm/redhat_transparent_hugepage/defrag`
 2. `echo no > /sys/kernel/mm/redhat_transparent_hugepage/khugepaged/defrag`
 3. Then issue each command in the shell so they are active immediately.

Note: These examples are applicable to Redhat 6.x settings. Check your system documentation for your platform.
- 8 Set the scheduler to deadline by adding the following line to `/etc/rc.local`
`echo deadline > /sys/block/sda/queue/scheduler`
Issue it on the command line so it is active immediately. This example assumes the drive is `sda`. Your drive may be different.

- 9 Set the blockdev for the drive on which vertica will run to 2048 by adding the following line to `/etc/rc.local` and run it in the shell so it is active immediately.
`/sbin/blockdev --setra 2048 /dev/sda2`

This example assumes the drive is `sda`. Your drive may be different.

The host is now ready for validation by the MC. If additional errors are found in the host during the validation steps, MC provides additional details about the specific error and steps on how to correct the problem.

Create a Private Key File

Before you can install a cluster, Management Console must be able to access the hosts on which you plan to install HP Vertica. MC uses password-less SSH to connect to the hosts and install HP Vertica software using a private key (pem) file.

If you already have a private key file, you can use it to connect to all potential hosts over SSH. In this case, you can skip creating the private key file and use your existing private key file.

Note: The private key file is required to complete the MC cluster installation wizard.

How private key files are created

Use the `/opt/vconsole/sbin/createRootPem` utility to create a private key file for the root user. The `createRootPem` utility has a single option, `-h`, which takes a comma-separated lists of host names or IP address for each node to include in the cluster. You must also include the MC host address if MC is installed outside the HP Vertica cluster.

After you run the command, you will be prompted for the root password for each of the hosts, as keys are generated and installed on each server.

When the process completes, each host has a new `/root/.ssh/authorized_keys2` file. If you installed MC outside of the cluster, the MC server has a file in `/root/.ssh` named `vid_rsa`.

Create a private key file

- 1 Log in on the server as root or as a user with sudo privileges.
- 2 Run `createRootPem` with the `-h` option and provide a list of hosts.
 - Example 1: Create a pem file on a 3-node cluster, where MC is already installed on one of those nodes:
`# createRootPem -h node01,node02,node03`
 - Example 2: Create a pem file on a 3-node cluster, where MC is installed on a dedicated server outside the cluster:
`# createRootPem -h mc01,node01,node02,node03`

Note: The order of the host arguments does not matter.

- 3 Enter the root password when prompted.
- 4 Copy the `vid_rsa` file to the local machine where you are going to run the browser session to connect to MC. The file is required by the cluster installation wizard, and you are prompted for the file from the wizard.

Use MC's Cluster Installation Wizard

MC's Cluster Installation Wizard guides you through the steps required to install an HP Vertica cluster on hosts that do not already have HP Vertica software installed.

Note: If you are using MC with the HP Vertica AML on Amazon Web Services, note that the Create Cluster and Import Cluster options are not supported.

Prerequisites

Before you proceed, make sure you:

- **Installed and configured MC** (page [58](#)).
- **Prepared the hosts** (page [51](#)) that you will include in the HP Vertica database cluster.
- **Created the private key (pem) file** and copied it to your local machine
- Obtained a copy of your HP Vertica license if you are installing the Enterprise Edition. If you are using the Community Edition, a license key is not required.
- Downloaded the HP Vertica server RPM (or DEB file).
- Have read/copy permissions on files stored on the local browser host that you will transfer to the host on which MC is installed.

Permissions on files you'll transfer to MC

On your local workstation, you must have at least read/write privileges on files you'll upload to MC through the Cluster Installation Wizard. These files include the HP Vertica server package, the license key (if needed), the private key file, and an optional CSV file of IP addresses.

Create a new HP Vertica cluster using MC

- 1 Connect to Management Console and log in as an MC administrator.
- 2 On MC's Home page, click the **Databases and Clusters** task.
- 3 Click the green plus sign and select **Create Cluster**.
- 4 The Create Cluster wizard opens. Provide the following information:
 1. Cluster name—A label for the cluster
 2. Vertica Admin User—The user that is created on each of the nodes when they are installed, typically 'dbadmin'. This user has access to HP Vertica and is also an OS user on the host.
 3. Password for the HP Vertica Admin User—The password you enter (required) is set for each node when MC installs HP Vertica.

Note: MC does not support an empty password for the administrative user.
 4. HP Vertica Admin Path—Storage location for catalog files, which defaults to /home/dbadmin unless you specified a different path during MC configuration (or later on MC's Settings page).

IMPORTANT: The Vertica Admin Path must be the same as the Linux database administrator's home directory. If you specify a path that is not the Linux dbadmin's home directory, MC returns an error.
- 5 Click **Next** and specify the private key file and host information:

1. Click **Browse** and navigate to the private key file (vid_rsa) that you created earlier.
Note: You can change the private key file at the beginning of the validation stage by clicking the name of the private key file in the bottom-left corner of the page. However, you cannot change the private key file after validation has begun unless the first host fails validation due to an SSH login error.
2. Include the host IP addresses. Here you have three options:
Specify later (but include number of nodes). This option allows you to specify the number of nodes, but not the specific IPs. You can specify the specific IPs before you validate hosts.
Import IP addresses from local file. You can specify the hosts in a CSV file using either IP addresses or host names.
Enter a range of IP addresses. You can specify a range of IPs to use for new nodes. For example 192.168.1.10 to 192.168.1.30. The range of IPs must be on the same or contiguous subnets.
- 6 Click **Next** and select the software and license:
 1. Vertica Software. If one or more HP Vertica packages have been uploaded, you can select one from the list; otherwise select **Upload a new local vertica binary file** and browse to an HP Vertica server file on your local system.
 2. Vertica License. Click **Browse** and navigate to a local copy of your HP Vertica license if you are installing the <ENT>. Community Edition versions require no license key
- 7 Click **Next**. The Create cluster page opens. If you did not specify the IP addresses, select each host icon and provide an IP address by entering the IP in the box and clicking **Apply** for each host you add.

The hosts are now ready for Host Validation and Cluster Creation.

Validate Hosts and Create the Cluster

Host validation is the process where the MC runs tests against each host in a **proposed cluster** (page 54).

You can validate hosts only after you have completed the cluster installation wizard. You must validate hosts before the MC can install HP Vertica on each host.

At any time during the validation process, but before you create the cluster, you can add and remove hosts by clicking the appropriate button in the upper left corner of the page on MC. A Create Cluster button appears when all hosts that appear in the node list are validated.

How to validate hosts

To validate one or more hosts:

- 1 Connect to Management Console and log in as an MC administrator.
- 2 On the MC Home page, click the **Databases and Clusters** task.
- 3 In the list of databases and clusters, select the cluster on which you have recently run the cluster installation wizard (**Creating...** appears under the cluster) and click **View**.

4 Validate one or several hosts:

- To validate a single host, click the host icon, then click **Validate Host**.
- To validate all hosts at the same time, click **All** in the Node List, then click **Validate Host**.
- To validate more than one host, but not all of them, Ctrl+click the host numbers in the node list, then click **Validate Host**.

5 Wait while validation proceeds.

The validation step takes several minutes to complete. The tests run in parallel for each host, so the number of hosts does not necessarily increase the amount of time it takes to validate all the hosts if you validate them at the same time. Hosts validation results in one of three possible states:

- Green check mark—The host is valid and can be included in the cluster.
- Orange triangle—The host can be added to the cluster, but warnings were generated. Click the tests in the host validation window to see details about the warnings.
- Red X—The host is not valid. Click the tests in the host validation window that have red X's to see details about the errors. You must correct the errors re-validate or remove the host before MC can create the cluster.

To remove an invalid host: Highlight the host icon or the IP address in the Node List and click **Remove Host**.

All hosts must be valid before you can create the cluster. Once all hosts are valid, a **Create Cluster** button appears near the top right corner of the page.

How to create the cluster

- 1 Click **Create Cluster** to install HP Vertica on each host and assemble the nodes into a cluster. The process, done in parallel, takes a few minutes as the software is copied to each host and installed.
- 2 Wait for the process to complete. When the **Success** dialog opens, you can do one of the following:
 - Optionally create a database on the new cluster at this time by clicking **Create Database**
 - Click **Done** to create the database at a later time

See **Creating a Database on a Cluster** for details on creating a database on the new cluster.

Create a Database on a Cluster

After you use the **MC Cluster Installation Wizard** (page [54](#)) to create an HP Vertica cluster, you can create a database on that cluster through the MC interface. You can create the database on all cluster nodes or on a subset of nodes.

If a database had been created using the Administration Tools on any of the nodes, MC detects (autodiscovers) that database and displays it on the Manage (Cluster Administration) page so you can import it into the MC interface and begin monitoring it.

MC allows only one database running on a cluster at a time, so you might need to stop a running database before you can create a new one.

The following procedure describes how to create a database on a cluster that you created using the MC **Cluster Installation Wizard** (page 54). To create a database on a cluster that you created by running the `install_vertica` script, see [Creating an Empty Database](#).

Create a database on a cluster

To create a new empty database on a new cluster:

- 1 If you are already on the **Databases and Clusters** page, skip to the next step; otherwise:
 - a) Connect to MC and sign in as an MC administrator.
 - b) On the Home page, click the **Databases and Clusters** task.
- 2 If no databases exist on the cluster, continue to the next step; otherwise:
 - a) If a database is running on the cluster on which you want to add a new database, select the database and click **Stop**.
 - b) Wait for the running database to have a status of *Stopped*.
- 3 Click the cluster on which you want to create the new database and click **Create Database**.
- 4 The Create Database wizard opens. Provide the following information:
 - Database name and password. See [Creating a Database Name and Password](#) for rules.
 - Optionally click **Advanced** to open the advanced settings and change the port and catalog, data, and temporary data paths. By default the MC application/web server port is 5450 and paths are `/home/dbadmin`, or whatever you defined for the paths when you ran the cluster creation wizard. Do not use the default agent port 5444 as a new setting for the MC application/web server port. See **MC Settings > Configuration** for port values.
- 5 Click **Continue**.
- 6 Select nodes to include in the database.

The Database Configuration window opens with the options you provided and a graphical representation of the nodes appears on the page. By default, all nodes are selected to be part of this database (denoted by a green check mark). You can optionally click each node and clear **Include host in new database** to exclude that node from the database. Excluded nodes are gray. If you change your mind, click the node and select the **Include** check box.
- 7 Click **Create** in the Database Configuration window to create the database on the nodes.

The creation process takes a few moments and then the database is started and a **Success** message appears.
- 8 Click **OK** to close the success message.

The Database Manager page opens and displays the database nodes. Nodes not included in the database are gray.

Installing and Configuring Management Console

This section describes how to install, configure, and upgrade Management Console (MC). If you need to back up your instance of MC, see *Backing Up MC* in the Administrator's Guide.

You can install MC before or after you install HP Vertica; however, consider installing HP Vertica and creating a database before you install MC. After you finish configuring MC, it automatically discovers your running database cluster, saving you the task of importing it manually.

Before You Install MC

Each version of HP Vertica Management Console (MC) is compatible only with the matching version of the HP Vertica server. For example, HP Vertica 6.1.2 server is supported with HP Vertica 6.1.2 MC only. Read the following documents for more information:

- Supported Platforms document, at <http://www.vertica.com/documentation> (<http://www.vertica.com/documentation>). The Supported Platforms document also lists supported browsers for MC.
- **Installation Overview and Checklist** (page [7](#)). Make sure you have everything ready for your HP Vertica configuration.
- **Before you Install HP Vertica** (page [12](#)). Read for required prerequisites for *all* HP Vertica configurations, including Management Console.

Driver requirements for Linux SuSe distributions

The MC (`vertica-console`) package contains the Oracle Implementation of Java 6 JRE and requires that you install the unixODBC driver manager on SuSe Linux platforms. unixODBC provides needed libraries `libodbc` and `libodbcinst`.

Port requirements

When you use MC to create a HP Vertica cluster, the **Create Cluster Wizard** (page [54](#)) uses SSH on its default port (22).

Port 5444 is the default agent port and must be available for MC-to-node and node-to-node communications.

Port 5450 is the default MC port and must be available for node-to-MC communications.

See **Ensure Ports are Available** (page [30](#)) for more information about port and firewall considerations.

Firewall considerations

Make sure that a firewall or iptables are not blocking communications between the cluster's database, Management Console, and MC's agents on each cluster node. See **Disable Firewalls** (page [30](#)).

IP address requirements

If you install MC on a server outside the HP Vertica cluster it will be monitoring, that server must be accessible to at least the public network interfaces on the cluster.

Disk space requirements

You can install MC on any node in the cluster, so there are no special disk requirements for MC—other than disk space you would normally allocate for your database cluster. See **Disk Space Requirements for HP Vertica** (page [17](#)).

Time synchronization and MC's self-signed certificate

When you connect to MC through a client browser, HP Vertica assigns each HTTPS request a self-signed certificate, which includes a timestamp. To increase security and protect against password replay attacks, the timestamp is valid for several seconds only, after which it expires.

To avoid being blocked out of MC, synchronize time on the hosts in your HP Vertica cluster, as well as on the MC host if it resides on a dedicated server. To recover from loss or lack of synchronization, resync system time and the Network Time Protocol. See **Set Up Time Synchronization** (page [90](#)) in the Installation Guide. If you want to generate your own certificates and keys for MC, see Generating Certifications and Keys for MC.

SSL requirements

The openssl package must be installed on your Linux environment so SSL can be set up during the MC configuration process. See SSL Prerequisites in the Administrator's Guide.

File permission requirements

On your local workstation, you must have at least read/write privileges on any files you plan to upload to MC through the **Cluster Installation Wizard** (page [54](#)). These files include the HP Vertica server package, the license key (if needed), the private key file, and an optional CSV file of IP addresses.

Monitor resolution

Management Console requires a minimum resolution of 1024 x 768, but HP recommends higher resolutions for optimal viewing.

Installing MC

You can install MC on any node you plan to include in the HP Vertica database cluster, as well as on its own, dedicated server outside the cluster.

Install MC

- 1 Download the MC package (`vertica-console-<current-version>.<architecture>`) from **myVertica portal** <http://my.vertica.com/> and save it to a location on the target server, such as `/tmp`.
- 2 On the target server, log in as root or a user with sudo privileges.

- 3 Change directory to the location where you saved the MC package.
- 4 Install MC using your local Linux distribution package management system (for example, rpm, yum, zipper, apt, dpkg).

The following command is a generic example for Red Hat 5:

```
# rpm -Uvh vertica-console-<current-version>.x86_64.RHEL5.rpm
```

The following command is a generic example for Debian 5 and Debian 6:

```
# dpkg -i vertica-console-<current-version>.deb
```

- 5 Open a browser and enter the IP address or host name of any cluster node, and include the default MC port 5450. For example, you'll enter one of:

```
https://xx.xx.xx.xx:5450/
```

```
https://hostname:5450/
```

Note: If you are installing MC on a server outside the HP Vertica cluster, you must specify the IP or hostname to that server, along with port 5450.

- 6 When the configuration wizard dialog box appears, proceed to **Configuring MC** (page [60](#)).

The files and scripts that MC needs to run are installed in `/opt/vconsole`, whereas HP Vertica is installed in `/opt/vertica`. Different locations prevent directory ownership/permission issues between Linux administrative users when you install MC on an HP Vertica node.

Configuring MC

After you **install MC** (page [59](#)), you need to configure it through a client browser connection. An MC configuration wizard walks you through creating the Linux MC super administrator account, storage locations, and other settings that MC needs to run. Information you provide during the configuration process is stored in the `/opt/vconsole/config/console.properties` file.

If you need to change settings after the configuration wizard ends, such as port assignments, you can do so later through Home > MC Settings page.

How to configure MC

- 1 Open a browser session.
- 2 Enter the IP address or host name of the server on which you installed MC (or any cluster node's IP/host name if you already installed HP Vertica), and include the default MC port 5450. For example, you'll enter one of:

```
https://xx.xx.xx.xx:5450/
```

```
https://hostname:5450/
```

- 3 Follow the configuration wizard.

About authentication for the MC super administrator

In the final step of the configuration process, you choose an authentication method for the MC super administrator. You can decide to have MC authenticate the MC super (in which case the process is complete), or you can choose LDAP.

If you choose LDAP, provide the following information for the newly-created MC super administrator:

- Corporate LDAP service host (IP address or host name)
- LDAP server running port (default 389)
- LDAP DN (distinguished name) for base search/lookup/authentication criteria
At a minimum, specify the dc (domain component) field. For example: `dc=vertica, dc=com` generates a unique identifier of the organization, like the corporate Web URL `vertica.com`
- Default search path for the organization unit (ou)
For example: `ou=sales, ou=engineering`
- Search attribute for the user name (uid), common name (cn), and so on
For example, `uid=jdoe, cn=Jane Doe`
- Binding DN and password for the MC super administrator.
In most cases, you provide the "Bind as administrator" fields, information used to establish the LDAP service connection for all LDAP operations, like search. Instead of using the administrator user name and password, the MC administrator could use his or her own LDAP credentials, as long as that user has search privileges.

If you choose bind anonymously

Unless you specifically configure the LDAP server to deny anonymous binds, the underlying LDAP protocol will not cause MC's **Configure Authentication** (page [60](#)) process to fail if you choose "Bind anonymously" for the MC administrator. Before you use anonymous bindings for LDAP authentication on MC, be sure that your LDAP server is configured to explicitly disable/enable this option. For more information, see the article on **Infusion Technology Solutions** <http://blog.infusiontechsolutions.com/disable-anonymous-access-to-openldap/>, as well as the **OpenLDAP documentation** <http://www.openldap.org/doc/admin24/access-control.html> on access control.

What happens next

Shortly after you click Finish, you should see a status in the browser; however, for several seconds you might see only an empty page. During this brief period, MC runs as the local user 'root' long enough to bind to port number 5450. Then MC switches to the MC super administrator account that you just created, restarts MC, and displays the MC login page.

Where to go next

If you are a new MC user and this is your first MC installation, you might want to familiarize yourself with MC design. See Management Console in the Concepts Guide.

If you'd rather use MC now, the following following topics in the Administrator's Guide should help get you started:

If you want to ...	See ...
Use the MC interface to install HP Vertica on a cluster of hosts	Creating an HP Vertica Cluster Using MC (page 51)
Create a new, empty HP Vertica database or import an	Managing Database Clusters on MC

existing HP Vertica database cluster into the MC interface	
Create new MC users and map them to one or more HP Vertica databases that you manage through the MC interface	Managing Users and Privileges (About MC Users and About MC privileges and roles)
Monitor MC and one or more MC-managed HP Vertica databases	Monitoring HP Vertica Using Management Console
Change default port assignments or upload a new HP Vertica license or SSL certificate	Managing MC Settings
Compare MC functionality to functionality that the Administration Tools provides	Administration Tools and Management Console

After You Install HP Vertica

The tasks described in this section are optional and are provided for your convenience. When you have completed this section, proceed to one of the following:

- Tutorial: Setting up an Example Database in the Getting Started Guide
- Configuring the Database in the Administrator's Guide

Install the License Key

If you did not supply the `-L` parameter during setup, or if you did not bypass the `-L` parameter for a **silent install** (page [48](#)), the first time you log in as the Database Administrator and run the HP Vertica Administration Tools or Management Console, HP Vertica requires you to install a license key.

Follow the instructions in Managing licenses in Administrator's Guide.

Verify Kernel and User Parameters Were Set

During installation, HP Vertica attempts to set various OS level parameters as shown in the Table below.

Note: The values listed below are the minimum recommended settings. The installer may not change values on your system if the value of the setting exceeds the value required by HP Vertica.

Setting / How to determine setting	Recommended Minimum Value	Comment
NOFILE <code>ulimit -n</code>	65536	Sets the maximum number of open files for the user. Change the setting in <code>/etc/security/limits.conf</code>
NPROC <code>ulimit -u</code>	Total RAM on the node in MB (for example on a 1GB machine, set to 1024)	Controls the maximum number of threads and processes for each user. Change the setting in <code>/etc/security/limits.conf</code>
fs.file-max <code>cat /proc/sys/fs/file-max</code>	65536	Sets the maximum number of file handles that the Linux kernel will allocate. Change the setting in <code>/etc/sysctl.conf</code>

vm.max_map_count <code>cat /proc/sys/vm/max_map_count</code>	Total RAM on the node in KB/16	The maximum number of memory map areas a process might have. This is set by taking the MemTotal value defined in <code>/proc/meminfo</code> and dividing by 16. Change the setting in <code>/etc/sysctl.conf</code>
readahead <code>blockdev --report</code> Look for the RA value. It is reported in KB.	2048	The readahead parameter for each partition. Change the setting by adding the following to <code>/etc/rc.local</code> <code>blockdev --setra 2048 /dev/sda2</code> Replace <code>/dev/sda2</code> with your partition(s). Do this for all HP Vertica partitions.
vm.min_free_kbytes <code>cat /proc/sys/vm/min_free_kbytes</code>	4096 or the value of the following formula; whichever is greater: <code>memtot=`grep MemTotal /proc/meminfo awk '{printf "%.0f", \$2}'` <code>echo "scale=0;sqrt (\$memtot*16)" bc</code></code>	Changes the page reclaim thresholds. When this number is increased the system starts reclaiming memory earlier, when its lowered it starts reclaiming memory later. Change the setting in <code>/etc/sysctl.conf</code>
fszize <code>ulimit -f</code>	unlimited	Sets the maximum file size to the OS limit. Change the setting in <code>/etc/security/limits.conf</code>

After installation is complete, check that the above parameters were correctly set. If for any reason, the installation script failed to do so, you might see warning messages in `adminTools-dbadmin.log`, similar to the following:

```
Feb 5 13:28:56 dbadmin@host01: /opt/vertica/bin/vertica --status -D
/vertica/dbs//amdaily/v_amdaily_node0001_catalog
Feb 5 13:28:56 dbadmin@nohost01: ['1', ["Not enough open file handles allowed (1024 available/32768
required); see 'ulimit -n'."]]
Feb 5 13:28:56 Error code 1 ["Not enough open file handles allowed (1024 available/32768 required);
see 'ulimit -n'."]
```

You can verify and fix the settings manually using the examples in ***Troubleshooting the HP Vertica Install*** (page [80](#)).

- ***Increase the Maximum Number of Files Open*** (page [88](#)) - use this procedure for any settings in `/etc/security/limits.conf` settings.

- **Increase the `max_map_count` Parameter** (page [89](#)) - use this procedure for any settings in `/etc/sysctl.conf`.

Optionally Install vsql Client Application on Non-cluster Hosts

You can use the HP Vertica vsql executable image on a non-cluster Linux host to connect to an HP Vertica database.

- On Red Hat 5.0 64-bit and SUSE 10/11 64-bit, you can install the client driver RPM, which includes the vsql executable. See [Installing the Client RPM on Red Hat and SUSE](#) for details.
- If the non-cluster host is running the same version of Linux as the cluster, copy the image file to the remote system. For example:

```
$ scp host01:/opt/vertica/bin/vsql .
$ ./vsql
```

- If the non-cluster host is running a different version of Linux than your cluster hosts, and that operating system is not Red Hat version 5 64-bit or SUSE 10/11 64-bit, you must install the HP Vertica server RPM in order to get vsql. Download the appropriate rpm package from the Download tab of the **myVertica portal** <http://my.vertica.com/> then log into the non-cluster host as root and install the rpm package using the command:

```
# rpm -Uvh filename
```

In the above command, *filename* is package you downloaded. Note that you do not have to run the `install_HP_Vertica` script on the non-cluster host in order to use vsql.

Notes

- Use the same command line options that you would on a cluster host.

You cannot run vsql on a Cygwin bash shell (Windows). Use ssh to connect to a cluster host, then run vsql. In release 5.1.5 vsql is also available for additional platforms. See [Installing the vsql client](#).

Install the HP Vertica-Ganglia Monitoring Package

The HP Vertica Analytics Platform is integrated with Ganglia, a web-based monitoring tool that lets you observe the status of an HP Vertica cluster and its running databases from your browser.

See [Monitoring HP Vertica Using Ganglia](#) in the Administrator's Guide.

Resolve SUSE Spread Configuration Issues

If a SUSE system does not have NTP running on startup, the HP Vertica installation could fail during spread configuration.

To resolve this issue, run the following command:

```
chkconfig --list ntp
```

The output resembles the following:

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

If the output does not look like the above example, run the following command:

```
chkconfig ntp on
```

SUSE users can now run `install_vertica` without spread configuration issues.

Changing the IP Addresses of an HP Vertica Cluster

This section describes how to change the IP addresses of the nodes in an HP Vertica cluster.

Note: This process requires that you stop the database on all nodes, then subsequently stop the database on individual nodes as you update IP addresses.

These instructions assume you will make the standard OS changes to change the IPs (for example, updating `/etc/hosts`) in Step 4 of this procedure. These instructions detail only the HP Vertica-specific IP changes. Consult the documentation for your particular OS platform for details on changing the IP address of the host.

To change the IP address of one or more nodes in a cluster:

- 1 Before changing the IP address on the Host, back up the following three files on all nodes:

- `/opt/vertica/config/admintools.conf`
- `/opt/vertica/config/vspread.conf`
- `/etc/sysconfig/spreadd`

- 2 Stop HP Vertica on all nodes.

- 3 As root, on each node, stop spread by using the following command:

```
/etc/init.d/spreadd stop
```

- 4 Change the IP addresses of the hosts as required by your operating system platform.

- 5 On each node edit `/opt/vertica/config/admintools.conf` and change the IPs as required.

You can use `sed` to change each IP in the file, for example to change `10.10.81.8` to `192.168.150.108` issue the command:

```
sed -i 's/10.10.81.8/192.168.150.108/g'  
/opt/vertica/config/admintools.conf
```

- 6 On each node edit `/opt/vertica/config/vspread.conf`:

1. Change the old IPs to the new IPs as required.

2. Locate the N number for each IP and change it to match the new IP. For example, if the old IP is 10.10.81.8, then the corresponding N number is N010010081008. N numbers consist of 3 digits for each IP number segment, padded with zeros when appropriate (10 becomes 010, 8 become 008, etc.). If the new IP address is 192.168.150.255, then the new corresponding N number is N192168250255.
- 7 On each node edit `/etc/sysconfig/spreadd` and change the N number to that node's new N number as you specified in `vspread.conf`.
- 8 As root, on each node start spread by using the following command:


```
/etc/init.d/spreadd start
```
- 9 Start the database.
- 10 Run `vsq`.
- 11 In `vsq`, issue the following query to verify the new IP has been updated:


```
select host_name from host_resources;
```

 You can also verify IPs with the following shell commands:
 - `cat /var/log/spreadd.log`
 - `admintools -t list_host`
 - `cat /etc/hosts`
- 12 Update the database to use the new IPs for reporting node status:
 1. In `vsq`, issue the command `select node_name, node_address from v_catalog.nodes;` to show you the current node names configured.
 - a) For each node in the result, change the hostname to the new IP address. **Note:** the node must be down to change the IP using the alter node command. You must bring the node down before altering the `NODE_NAME` property for that node:
 - Bring the node down that you are going to update (don't bring down the node from which you are using `vsq`!). You can bring down the node from the initiator node using `admintools`. For example:


```
admintools -t stop_node -s 192.168.150.255
```
 - Update the IP address by issuing the command: `alter node NODE_NAME is hostname 'new.ip.address' ;` where `NODE_NAME` is the node_name and new.ip.address is the new IP address of that node.
 - Bring the node back up, check it's status in the nodes table (`select node_name, node_state from nodes;`), and wait for that node's status to be `UP`. You can use `admintools` from the initiator node to restart the node. You must provide the database name. For example:


```
admintools -t restart_node -s 192.168.150.255 -d VMart
```
 - Repeat the process for the next node in the result. After all nodes have been updated except the node from which you are using `vsq`, log out of `vsq`, then log into `vsq` from another node and update the IP address for the node from which you were previously using `vsq`.

See Also

Configure Hostname Resolution (page [27](#))

Install HP Vertica Documentation

The latest documentation for your HP Vertica release is available at <http://www.vertica.com/documentation> (<http://www.vertica.com/documentation>). After you install HP Vertica, install the documentation on your database server and client systems.

Installing the HP Vertica documentation set on database server systems

To install the documentation on the server host:

- 1 Log in as root or sudo on the target server:

```
# su - root
```
- 2 Open a Web browser and go to <http://www.vertica.com/documentation> (<http://www.vertica.com/documentation>).
- 3 Scroll down and save the HP Vertica documentation package (.tar.gz or .zip) to your system; for example, to /tmp.
- 4 Change directory to the location of the package:

```
# cd /tmp
```
- 5 Extract the contents:

```
# tar xzf vertica-user-doc-6.1.x.x.tar.gz
```

The extraction process creates a /vertica-user-doc-6.1.x.x/ folder with the following directories and contents:

 - HTML/index.htm contains the index.htm file for full HP documentation set
 - JDBC/index.htm provides access to the JavaDoc HTML files
 - PDF/book-name.pdf contains a PDF for each book

Installing the HP Vertica documentation set on on database client systems

To install a local copy of the documentation on a client system:

- 1 Open a Web browser and go to <http://www.vertica.com/documentation> (<http://www.vertica.com/documentation>).
- 2 Scroll down and save the HP Vertica documentation package (.tar.gz or .zip) to your system; for example, to /tmp.
- 3 Change directory to the location of the package:

```
$ cd /tmp
```
- 4 Extract the contents:

```
$ tar xzf vertica-user-doc-6.1.x.x.tar.gz
```

The extraction process creates a /vertica-user-doc-6.1.x.x/ folder with the following directories and contents:

 - HTML/index.htm contains the index.htm file for full HP documentation set
 - JDBC/index.htm provides access to the JavaDoc HTML files
 - PDF/book-name.pdf contains a PDF for each book

Note: If you have a browser or Adobe Acrobat Reader installed on the cluster host or on each client, you can also access the documentation directly from the **Product Documentation page** (<http://www.vertica.com/documentation>), but you need a V-Zone login. All cross-references within the HTML documentation are relative, so there is no location dependency. This is the method that HP recommends.

Get Started!

HP Vertica lets you choose between instant gratification and a more detailed path in setting up your example database. Both processes, described in the Getting Started Guide, are simple, and both let you start using your database immediately—literally within minutes.

- If you can't wait to get started, read about one-step installation script in Installing the Example Database.
- If you prefer a more thorough, but equally-simple example, see the Tutorial.

Installing Client Drivers

After you install HP Vertica, install drivers on the client systems from which you plan to access your databases. HP supplies drivers for ADO.NET, JDBC, ODBC Perl, and Python. For instructions on installing these drivers, see Client driver install procedures in the Programmer's Guide.

Extend Your Installation Using HP Vertica Packages

You can extend the functionality of your HP Vertica installation in a number of different ways. You can create your own SQL functions, which allow you to execute complex queries and combine HP Vertica built-in functions. You can also create your own user-defined functions (UDFs), which are libraries of functions that you develop in C++ to perform analytic operations that are difficult to perform with SQL. And, you can take advantage of HP Vertica *packages* to install additional SQL functions to your system. An example of a package is the HP Vertica Geospatial package, which allows you to query and analyze geographic/spatial data.

Installing packages

HP Vertica packages are optional; you can decide whether or not you want to install and use them. After you install HP Vertica, you'll find all of the directories and files necessary for package installation within the `/opt/vertica/packages/<package-name>` directory.

Keep in mind that, if you decide to make any changes to the SQL functions in any package, you should be sure to save a backup of your changes in another directory, so that your changes will be preserved in the event of an upgrade. HP Vertica is not responsible for migrating any changes you make to package functions to newer versions.

Directory structure for HP Vertica packages

At the top level of each package directory, you'll find these files:

- `readme.txt`—contains information on how to install the package
- `install.sh`—use this script to install the package

Each package directory might also contain a number of subdirectories. Note that each of these subdirectories is created only when the package requires the specified files.

- `/bin`—contains procedures needed for the package.
- `/lib`—contains shared libraries (.so files) needed for the package.
- `/src`—contains .sql or other source files needed for the package.
- `/examples`—contains example files (.demo.sql files) that illustrate common use cases.

To install an HP Vertica package:

Run the `install.sh` script that appears in the `/opt/vertica/packages/` directory for the specific package.

Note: If you choose to install an HP Vertica package in a directory other than the default, be sure to set the appropriate environment variable to reflect the correct directory.

Upgrading HP Vertica

Upgrading HP Vertica to a New Version

The process of upgrading HP Vertica is similar to installing it.

Important! When upgrading from 5.x to 6.0, due to a change in how transaction catalog storage works in HP Vertica 6.0 and later, the amount of space that the transaction catalog takes up can increase significantly during and after the upgrade. Verify that you have at least 4x the size of the Catalog folder in the catalog free (in addition to normal free space requirements) on your nodes prior to upgrading.

To determine the amount of space the Catalog folder is using, run `du -h` on the Catalog folder.

Note: do not run `du -h` on the entire catalog. Run it specifically on the Catalog folder in the catalog.

For example:

```
[dbadmin@localhost ~]$ du -h
/home/dbadmin/db/v_db_node0001_catalog/Catalog/
```

Upgrading HP Vertica

Follow these steps to upgrade your database. Note that upgrades are incremental and must follow one of the following upgrade paths:

- HP Vertica 3.5 to 4.0
- HP Vertica 4.0 to 4.1
- HP Vertica 4.1 to 5.0
- HP Vertica 4.1 to 5.1
- HP Vertica 5.0 to 5.1
- HP Vertica 5.0 to 6.0
- HP Vertica 5.1 to 6.0
- HP Vertica 6.0 to 6.1

IMPORTANT NOTE: HP strongly recommends that you follow the upgrade paths. Be sure to read the New Features and New Features for each version you skip. The HP Vertica documentation is available in the rpm, as well as at <http://www.vertica.com/documentation> (<http://www.vertica.com/documentation>) (which also provides access to previous versions of the documentation).

- 1 Back up your existing database. This is a precautionary measure so that you can restore from the backup if the upgrade is unsuccessful.

Note: Release 5.1 introduced a new backup utility, `vbr.py`. This utility replaced both the `backup.sh` and `restore.sh` scripts, making both obsolete. Any backups created with `backup.sh` are incompatible with backups created with `vbr.py`. HP Vertica recommends that you use the current utility `vbr.py` as soon as possible after successfully upgrading to Release 5.1 or later. Documentation for the 5.0 scripts remained in the 5.1 documentation. However, the topics were marked Obsolete in that version and subsequently removed from the documentation.

- 2 If you are using Ganglia to monitor your database, disable any crontab entry you created for it. Otherwise, it could run during the upgrade and interfere with the installer. See [Monitoring HP Vertica Using Ganglia](#) for more information about Ganglia, and [Add a cron job](#) for instructions on editing the crontab.
- 3 Stop the database using admintools if it is not already stopped.
- 4 On each host that you have an additional package installed, such as the R Language Pack, uninstall the package. For example: `rpm -e vertica-R-lang`.
- 5 On the administration host, install the new HP Vertica package. See **Download and install the HP Vertica install package** (page [34](#)). For example:

```
rpm syntax:
# rpm -Uvh /home/dbadmin/vertica-x86_64.RHEL5.rpm

deb syntax:
# dpkg -i /home/dbadmin/vertica-x86_64.RHEL5.rpm
```

Note: If you fail to install the rpm or deb prior to running the next step, then `update_vertica` fails with an error due to the conflict between the version of the `update_vertica` script and the version of the rpm argument.

- 6 As root or sudo, run `update_vertica`. **Use the same options that you used when you last installed or upgraded the database, except for the `-s host_list` parameter.** Use the `-S default` parameter if spread is running on the same subnet as HP Vertica (default behavior). Otherwise use the `-S` parameter with the same options you originally used when installing HP Vertica. If you forgot the options that were last used, open `/opt/vertica/config/admintools.conf` in a text editor and find the line that starts with `install_opts`. This line details each option.

For example:

```
# /opt/vertica/sbin/update_vertica -S default -r
/home/dbadmin/vertica-x86_64.RHEL5.rpm
```

Note: the rpm file must be readable by the dbadmin user when upgrading. Some upgrade scripts are run as the dbadmin user, and that user must be able to read the rpm file. Stop all databases before you use the `-S` parameter. If the databases are not stopped then the `-S` parameter causes the database to restart.

- 7 Start the database. The startup scripts analyze the database and perform any necessary data and catalog updates for the new version.
- 8 Stop the database and perform another backup. Note that you can use incremental backups as you upgrade, but when moving from Version 5.0 and earlier version to Version 5.1 and later versions the backup process changes from using `backup.sh` to using `vbr.py`. You cannot use an incremental backup with different version of the backup script. Create a full backup the first time you move to using `vbr.py`.
- 9 Continue along the upgrade path and perform these same steps for each version in your upgrade path.
- 10 After you have upgraded to the latest version of the server, install any additional packs you had previously installed. See the pack install/upgrade instructions for details on upgrading the packs. For R, see [Installing/Upgrading the R Language Pack for HP Vertica](#).

Additional notes

- Downgrade installations are not supported.
- HP Vertica includes a DBA UNIX group for tighter control over filesystem access in the `/opt/vertica/` directories.

During upgrade, a `verticadba` group is created, and existing users are added to the group with permissions set to `775`. This setting grants full privileges to the DBA 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.

See also

Run the install script (page [35](#))

Upgrading Your HP Vertica License

To upgrade from the Community Edition license, obtain an evaluation or Enterprise Edition license from HP. For information on applying your new license, see *Installing or Upgrading a License Key* in the Administrator's Guide.

Upgrading MC

If you are moving from MC 6.1.1 to MC 6.1.2, you can install MC on any HP Vertica cluster node. This scenario requires a fresh install because HP does not provide scripts to migrate metadata (MC users and settings) established in earlier releases from your existing server to the cluster node. See *Installing and Configuring Management Console* (page [58](#)).

After you install and configure MC, you will need to recreate MC users you'd created for your 6.1.1 MC instance, if any, and apply previous MC settings to the new MC version.

Tip: You can export MC-managed database messages and user activity to a location on the existing server. While you can't import this data, using the exported files as a reference could help make metadata recreation easier. See *Exporting MC-managed Database Messages and Exporting the Audit Log*.

If you want to keep MC on the existing server

If you want to keep MC on the same server (such as on the dedicated server that had been required in previous MC releases), your MC metadata is retained when you run the `vertica-console` installation script.

Before you upgrade MC on the same server

- 1 Log in as root or a user with sudo privileges on the server where MC is already installed.
- 2 Back up MC to preserve configuration metadata. See *Backing Up MC*.
- 3 Open a terminal window and shut down the MC process using the following command:

```
# /etc/init.d/vertica-consolled stop
```

Upgrade MC on the same server

- 1 Download the MC package
(`vertica-console-<current-version>.<architecture>`) from **myVertica portal** <http://my.vertica.com/> and save it to a location on the target server, such as `/tmp`.
- 2 On the target server, log in as root or a user with sudo privileges.
- 3 Change directory to the location where you saved the MC package.
- 4 Install MC using your local Linux distribution package management system (for example, rpm, yum, zipper, apt, dpkg).

The following command is a generic example for Red Hat 5:

```
# rpm -Uvh vertica-console-<current-version>.x86_64.RHEL5.rpm
```

The following command is a generic example for Debian 5 and Debian 6:

```
# dpkg -i vertica-console-<current-version>.deb
```

- 5 Open a browser and enter the IP address or host name of any cluster node, and include the default MC port 5450. For example, you'll enter one of:

`https://xx.xx.xx.xx:5450/`

`https://hostname:5450/`

Note: If you are installing MC on a server outside the HP Vertica cluster, you must specify the IP or hostname to that server, along with port 5450.

- 6 When the configuration wizard dialog box appears, proceed to **Configuring MC** (page [60](#)).

Uninstalling HP Vertica

Uninstalling HP Vertica

To uninstall HP Vertica:

- 1 For each host in the cluster, do the following:
 1. Choose a host machine and log in as root (or log in as another user and switch to root).

```
$ su - root
password: <root-password>
#
```

2. Find the name of the package that is installed:

```
# rpm -qa | grep vertica
```

3. Remove the package:

```
# rpm -e package
```

Note: If you want to delete the configuration file used with your installation, you can choose to delete the `/opt/vertica/` directory and all subdirectories using this command: `# rm -rf /opt/vertica/`

- 2 For each client system, do the following:

1. Delete the JDBC driver jar file.
2. Delete ODBC driver data source names.
3. Delete the ODBC driver software. On Windows, go to Start > Control Panel > Add or Remove Programs and look for HP Vertica. Click **Remove**.

Uninstalling MC

The uninstall command shuts down Management Console and removes most of the files that MC installation script installed.

Uninstall MC

- 1 Log in to the target server as root:
- 2 Look for previously-installed versions of MC and note the version:

```
# rpm -qa | grep vert
```

For example, the following is similar to output you will see:

```
vertica-console-6.1.2-0.x86_64
```

- 3 Run the erase command:

```
# rpm -e <vertica-console>
```

- 4 Check for running processes and shut them down, including spread.

1. Check for running processes:

```
# ps ax | grep opt
```

2. If you see nothing running, proceed to step 5. If output similar to the following appears, proceed to step 4c:

```
3170 ? S<s 0:00 /opt/vertica/spread/sbin/spread -n N127000000001 -c
```

```
          /opt/vertica/config/vspread.conf
3292 ? S<s 0:00 /bin/sh /opt/vertica/agent/agent.sh
          /opt/vertica/config/users/mcadmin/agent.conf
3334 ? S<l 0:33 /opt/vertica/oss/python/bin/python ./simply_fast.py
```

3. Stop spread:

```
# /etc/init.d/spreadd stop
```

4. Issue the kill command on the remaining running processes, using their process ID (PID); for example, based on the sample output in step 4a:

```
# kill -9 3292
# kill -9 3334
```

5 Manually delete the /opt/vconsole directory:

```
# rm -rf /opt/vconsole
```

If you want to reinstall MC

To re-install MC, see *Installing and Configuring Management Console* (page [58](#)).

Troubleshooting the HP Vertica Install

The topics described in this section are performed automatically by the `install_vertica` script and are described in *Installing HP Vertica* (page [33](#)). If you did not encounter any installation problems, proceed to the Administrator's Guide for instructions on how to configure and operate a database.

Validation Scripts

HP Vertica provides several validation utilities that can be used prior to deploying HP Vertica to help determine if your hosts and network can properly handle the processing and network traffic required by HP Vertica. These utilities can also be used if you are encountering performance issues and need to troubleshoot the issue.

After you install the HP Vertica RPM, you have access to the following scripts in `/opt/vertica/bin`:

- ***vcupperf*** - a CPU performance test used to verify your CPU performance.
- ***vioperf*** - an Input/Output test used to verify the speed and consistency of your hard drives.
- ***vnetperf*** - a Network test used to test the latency and throughput of your network between hosts.

These utilities can be run at any time, but are well suited to use before running the `install_vertica` script.

vcupperf

The `vcupperf` utility measures your server's CPU processing speed and compares it against benchmarks for common server CPUs. The utility performs a CPU test and measures the time it takes to complete the test. The lower the number scored on the test, the better the performance of the CPU.

The `vcupperf` utility also checks the high and low load times to determine if CPU throttling is enabled. If a server's low-load computation time is significantly longer than the high-load computation time, CPU throttling may be enabled. CPU throttling is a power-saving feature. However, CPU throttling can reduce the performance of your server. HP Vertica recommends disabling CPU throttling to enhance server performance.

Syntax

```
vcupperf [-q]
```

Options

`-q`: run in quiet mode. Quiet mode displays only the CPU Time, Real Time, and high and low load times.

Returns

- CPU Time: the amount of time it took the CPU to run the test.
- Real Time: the total time for the test to execute.
- High load time: The amount of time to run the load test while simulating a high CPU load.
- Low load time: The amount of time to run the load test while simulating a low CPU load.

Example

The following example shows a CPU that is running slightly slower than the expected time on a Xeon 5670 CPU that has CPU throttling enabled.

```
[root@docb04 bin]# /opt/vertica/bin/vcpuperf
Compiled with: 4.1.2 20080704 (Red Hat 4.1.2-52)
Expected time on Core 2, 2.53GHz: ~9.5s
Expected time on Nehalem, 2.67GHz: ~9.0s
Expected time on Xeon 5670, 2.93GHz: ~8.0s
```

```
This machine's time:
  CPU Time: 8.540000s
  Real Time:8.710000s
```

Some machines automatically throttle the CPU to save power.

This test can be done in <100 microseconds (60-70 on Xeon 5670, 2.93GHz).

Low load times much larger than 100-200us or much larger than the corresponding high load time indicate low-load throttling, which can adversely affect small query / concurrent performance.

```
This machine's high load time: 67 microseconds.
This machine's low load time: 208 microseconds.
```

vioperf

The `vioperf` utility quickly tests the performance of your host's input and output subsystem. The utility performs the following tests:

- sequential write
- sequential rewrite
- sequential read
- skip read (read non-contiguous data blocks)

The utility verifies that the host reads the same bytes that it wrote and prints its output to STDOUT. The utility also logs the output to a JSON formatted file.

Minimum and Recommended IO Performance

- The minimum required I/O is 20 MB/s read/write per physical processor core on each node, in full duplex (reading and writing) simultaneously, concurrently on all nodes of the cluster.
- The recommended I/O is 40 MB/s per physical core on each node.

For example, the I/O rate for a node with 2 hyper-threaded six-core CPUs (12 physical cores) is 240 MB/s required minimum, 480 MB/s recommended.

Syntax

```
vioperf [--help] [--duration=<INTERVAL>] [--log-interval=<INTERVAL>]
[--log-file=<FILE>] [--condense-log] [<DIR>*]
```

Options

<code>--help</code>	Prints a help message and exits.
<code>--duration</code>	The length of time <code>vioprobe</code> runs performance tests. The default is 5 minutes. Specify the interval in seconds, minutes, or hours with any of these suffixes: <ul style="list-style-type: none">▪ Seconds: s, sec, secs, second, seconds. Example: <code>--duration=60sec</code>▪ Minutes: m, min, mins, minute, minutes. Example: <code>--duration=10min</code>▪ Hours: h, hr, hrs, hour, hours. Example: <code>--duration=1hrs</code>
<code>--log-interval</code>	The interval at which the log file reports summary information. The default interval is 10 seconds. This option uses the same interval notation as <code>--duration</code> .
<code>--log-file</code>	The path and name where log file contents are written, in JSON. If not specified, then <code>vioperf</code> creates a file named <code>results</code> <i>date-time</i> .JSON in the current directory.
<code>--condense-log</code>	Directs <code>vioperf</code> to write the log file contents in condensed format, one JSON entry per line, rather than as indented JSON syntax.
<code><DIR></code>	Zero or more directories to test. If you do not specify a directory, <code>vioperf</code> tests the current directory. To test the performance of each disk, specify different directories mounted on different disks.

Returns

The utility returns the following information:

test	The test being run (Write, ReWrite, Read, or Skip Read)
directory	The directory in which the test is being run.
counter name	The counter type of the test being run. Can be either MB/s or Seeks per second.
counter value	The value of the counter in MB/s or Seeks per second across all threads. This measurement represents the bandwidth at the exact time of measurement. Contrast with counter value (avg).
counter value (avg)	The average amount of data in MB/s, or the average number of Seeks per second, for the test being run in the duration specified with <code>--log-interval</code> . The default interval is 10 seconds. The <code>counter value (avg)</code> is the average bandwidth since the last log message, across all threads.
thread count	The number of threads used to run the test.
%CPU	The available CPU percentage used during this test.

%IO Wait	The CPU percentage in I/O Wait state during this test. I/O wait state is the time working processes are blocked while waiting for I/O operations to complete.
elapsed time	The amount of time taken for a particular test. If you run the test multiple times, elapsed time increases the next time the test is run.
remaining time	The time remaining until the next test. Based on the <code>--duration</code> option, each of the tests is run at least once. If the test set is run multiple times, then remaining time is how much longer the test will run. The remaining time value is cumulative. Its total is added to elapsed time each time the same test is run again.

Example

Invoking `vioperf` from a terminal outputs the following message and sample results:

```
[dbadmin@node01 ~]$ /opt/vertica/bin/vioperf --duration=60s
The minimum required I/O is 20 MB/s read and write per physical processor core on each node, in full
duplex i.e. reading and writing at this rate simultaneously, concurrently on all nodes of the cluster.
The recommended I/O is 40 MB/s per physical core on each node. For example, the I/O rate for a server
node with 2 hyper-threaded six-core CPUs is 240 MB/s required minimum, 480 MB/s recommended.

Using direct io (buffer size=1048576, alignment=512) for directory
"/home/dbadmin/VMart/v_vmart_node0001_catalog"
```

test	directory	counter name	counter value
counter value (10 sec avg)	thread count	%CPU	%IO Wait
(s)	elapsed time (s)	remaining time	

Write	/home/dbadmin/VMart/v_vmart_node0001_catalog	MB/s	27
27	1	2	97
Write	/home/dbadmin/VMart/v_vmart_node0001_catalog	MB/s	24
16	1	3	96
ReWrite	/home/dbadmin/VMart/v_vmart_node0001_catalog	MB/s	8
8	1	2	96
ReWrite	/home/dbadmin/VMart/v_vmart_node0001_catalog	MB/s	9
12	1	3	95
Read	/home/dbadmin/VMart/v_vmart_node0001_catalog	MB/s	22
22	1	4	94
Read	/home/dbadmin/VMart/v_vmart_node0001_catalog	MB/s	23
27	1	6	91
SkipRead	/home/dbadmin/VMart/v_vmart_node0001_catalog	seeks/s	300
300	1	0	99
SkipRead	/home/dbadmin/VMart/v_vmart_node0001_catalog	seeks/s	308
324	1	0	99

vnetperf

The `vnetperf` utility allows you to measure the network performance of your hosts. It can measure network latency and the throughput for both the TCP and UDP protocols.

IMPORTANT! This utility introduces a high network load and must not be used on a running HP Vertica cluster or database performance is degraded.

Using this utility you can detect:

- if throughput is low for all hosts or a particular host,
- if latency is high for all hosts or a particular host,
- bottlenecks between one or more hosts or subnets,
- too low a limit in the number of TCP connections that can be established simultaneously,
- and if there is a high rate of packet loss on the network.

The latency test measures the latency from the host running the script to the other hosts. Any host that has a particularly high latency should be investigated further.

The throughput tests measure both UDP and TCP throughput. You can specify a rate limit in MB/s to use for these tests, or allow the utility to use a range of throughputs to be used.

Recommended Network Performance

- The maximum recommended RTT (round-trip time) latency is 1000 microseconds. The ideal RTT latency is 200 microseconds or less. HP Vertica recommends that clock skew be kept to under 1 second.
- The minimum recommended throughput is 100MB/s. Ideal throughput is 800 MB/s or more.
Note: UDP numbers may be lower, multiple network switches may reduce performance results.

Syntax

```
vnetperf [options] [tests]
```

Options

<code>--condense</code>	Condense the log into one JSON entry per line, instead of indented JSON syntax.
<code>--collect-logs</code>	Collect the test log files from each host.
<code>--datarate <i>rate</i></code>	Limit the throughput to this rate in MB/s. A rate of 0 loops the tests through several different rates. The default is 0.
<code>--duration <i>seconds</i></code>	The time limit for each test to run in seconds. The default is 1.
<code>--hosts <i>host1,host2,...</i></code>	A comma-separated list of hosts on which to run the tests. Do not use spaces between the comma's and the host names.
<code>--hosts <i>file</i></code>	A hosts file that specifies the hosts on which to run the tests. If the <code>--hosts</code> argument is not used, then the utility attempts to access <code>admintools</code> and determine the hosts in the cluster.
<code>--identity-file <i>file</i></code>	If using passwordless SSH/SCP access between the hosts, then specify the key file used to gain access to the hosts.
<code>--ignore-bad-hosts</code>	If set, run the tests on the reachable hosts even if some hosts are not reachable. If not set, and a host is unreachable, then no tests are run on any hosts.
<code>--log-dir <i>directory</i></code>	If <code>--collect-logs</code> is set, the directory in which to place the collected logs. The default directory is named <code>logs.netperf.<timestamp></code>
<code>--log-level <i>LEVEL</i></code>	The log level to use. Possible values are: INFO, ERROR, DEBUG, and WARN. The default is WARN.

<code>--list-tests</code>	Lists the tests that can be run by this utility.
<code>--output-file <i>file</i></code>	The file that JSON results are written to. The default is results.<timestamp>.json.
<code>--ports <i>port1,port2,port3</i></code>	The port numbers to use. If only one is specified then the next two numbers in sequence are also used. The default ports are 14159,14160, 14161.
<code>--scp-options '<i>options</i>'</code>	Using this argument, you can specify one or more standard SCP command line arguments enclosed in single quotes. SCP is used to copy test binaries over to the target hosts.
<code>--ssh-options '<i>options</i>'</code>	Using this argument, you can specify one or more standard SSH command line arguments enclose in single quotes. SSH is used to issue test commands on the target hosts.
<code>--vertica-install <i>directory</i></code>	If specified, then the utility assumes HP Vertica is installed on each of the hosts and to use the test binaries on the target system rather than copying them over using SCP.

Tests

Note: If the tests argument is omitted then all tests are run.

latency	Test the latency to each of the hosts.
tcp-throughput	Test the TCP throughput amongst the hosts.
udp-throughput	Test the UDP throughput amongst the hosts.

Returns

For each host it returns the following:

Latency test returns:

- The Round Trip Time (rtt) latency for each host in milliseconds.
- Clock Skew = the difference in time shown by the clock on the target host relative to the host running the utility.

UDP and TCP throughput tests return:

- The date/time and test name.
- The rate limit in MB/s.
- The node being tested.
- Sent and Received data in MB/s and bytes.
- The duration of the test in seconds.

Example

```
/opt/vertica/bin/vnetperf --condense -hosts 10.20.100.66,10.20.100.67
--identity-file '/root/.ssh/vid_rsa'
```

Enable Port 5433

Ensure that 5433 is TCP-enabled on all hosts for connections inside and outside the database. This includes making sure that all of the following are correctly configured or disabled.

- Firewalls — not recommended for database hosts
- SELinux (Security-Enhanced Linux)
- iptables

If it is consistent with your security policy, you can disable the Linux firewall and SELinux by running the command:

```
system-config-securitylevel
```

For more information about SELinux

- *NSA's Page on SELinux* (<http://www.nsa.gov/research/selinux/index.shtml>)
- *Red Hat's Page on SELinux*
(http://www.redhat.com/promo/security/?sc_cid=MSI_GSEL_25&s_kwcid=selinux|731967321)

For more information about iptables

- *The netfilter.org "iptables" project*
(<http://www.netfilter.org/projects/iptables/index.html>)
- *Iptables Tutorial 1.2.1* (<http://www.frozentux.net/documents/iptables-tutorial/>)

Enable Secure Shell (SSH) Logins

The administrative account must be able to use Secure Shell (SSH) to log in (ssh) to all hosts without specifying a password. The shell script `install_vertica` does this automatically. This section describes how to do it manually if necessary.

- 1 If you do not already have SSH installed on all hosts, log in as root on each host and install it now. You can download a free version of the SSH connectivity tools from **OpenSSH** (<http://www.openssh.com/>).
- 2 Log in to the HP Vertica administrator account (dbadmin in this example).
- 3 Make your home directory (~) writable only by yourself. Choose one of:

```
$ chmod 700 ~
```

or

```
$ chmod 755 ~
```

where:

700 includes

400 read by owner

200 write by owner

100 execute by owner

755 includes

400 read by owner

200 write by owner

100 execute by owner

```
040 read by group
010 execute by group
004 read by anybody (other)
001 execute by anybody
```

4 Change to your home directory:

```
$ cd ~
```

5 Generate a private key/ public key pair:

```
$ ssh-keygen -t rsa
Generating public/private rsa key pair.
Enter file in which to save the key (/home/dbadmin/.ssh/id_rsa):
Created directory '/home/dbadmin/.ssh'.
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /home/dbadmin/.ssh/id_rsa.
Your public key has been saved in /home/dbadmin/.ssh/id_rsa.pub.
```

6 Make your .ssh directory readable and writable only by yourself:

```
$ chmod 700 ~/.ssh
```

7 Change to the .ssh directory:

```
$ cd ~/.ssh
```

8 Copy the file id_rsa.pub onto the file authorized_keys2.

```
$ cp id_rsa.pub authorized_keys2
```

9 Make the files in your .ssh directory readable and writable only by yourself:

```
$ chmod 600 ~/.ssh/*
```

10 For each cluster host:

```
$ scp -r ~/.ssh <host>:.
```

11 Connect to each cluster host. The first time you ssh to a new remote machine, you could get a message similar to the following:

```
$ ssh dev0
Warning: Permanently added 'dev0,192.168.1.92' (RSA) to the list of known hosts.
```

This message appears only the first time you ssh to a particular remote host.

See also

The **OpenSSH** (<http://www.openssh.com/>) Web site

Increase Free Memory

Note: This step is automatically done by the installer. See **Verify that Kernel and User Parameters Were Set** (page 63) for details on the setting. You only need to perform this step if the automatic setting did not work.

To increase free memory perform the following steps:

- 1 Determine how much free memory should be available. Run the following on the command line:

```
memtot=`grep MemTotal /proc/meminfo | awk '{printf "%.0f",$2}'`
echo "scale=0;sqrt($memtot*16)" | bc
```

- 2 If the number is greater than 4096, then use that number as the *VALUE* on the following lines. Otherwise use 4096 as the *VALUE*.

- 3 Add the following line to `/etc/sysctl.conf` (the kernel sysctl configuration file):

```
vm.min_free_kbytes = VALUE
```

For example:

```
vm.min_free_kbytes = 45056
```

- 4 At the root command line prompt, enter:

```
# echo VALUE > /proc/sys/vm/min_free_kbytes
```

For example:

```
# echo 45056 > /proc/sys/vm/min_free_kbytes
```

Increase Maximum Number of Files Open

Note: This step is automatically done by the installer. See *Verify that Kernel and User Parameters Were Set* (page 63) for details on the setting. You only need to perform this step if the automatic setting did not work.

HP Vertica requires a large number of file handles. To avoid "too many files open" errors:

- 1 Check the current limit:

```
$ ulimit -n
```

- 2 If necessary, permanently increase this limit for all users by editing the file `/etc/security/limits.conf` as root, and adding the lines:

```
* soft nofile <value>
* hard nofile <value>
```

<value> = number of MB of RAM, or 65536, whichever is higher.

Note: The leading asterisk means that the setting applies to all users.

- 3 Set `file-max` to the number of MB of RAM, or 65536, whichever is higher.

If you see many error messages about running out of file handles, try increasing this value; for example:

```
echo 100000 > /proc/sys/fs/file-max
```

For further details, see the `file-max` section on the *Linux man proc* page <http://linux.die.net/man/5/proc>.

- 4 Log out and log in again. The new shell has the new `ulimit` setting.

If you used sudo to install HP Vertica:

- 1 View the `/etc/security/limits.conf` on each node in the cluster and verify that HP Vertica inserted the following lines at the end of the file:

```
* soft nofile <value>
```



```
* hard nofile <value>
```

where <value> = number of MB of RAM, or 65536, whichever is higher.

Note: The leading asterisk means that the setting applies to all users.

- 2 Log in as root and restart the system to ensure the settings are implemented:


```
# etc/init.d/sshd restart
```

 If restart is not an option, stop and start sshd:


```
# etc/init.d/sshd stop; etc/init.d/sshd start
```
- 3 Verify the number of open files is 65536 (or 1024 on some operating systems):


```
# ulimit -n
65536
```
- 4 Log in as the dbadmin user:


```
# su dbadmin
```
- 5 SSH from one node to another; for example, from host01:


```
# ssh host2
```
- 6 Check the current file limit:


```
# ulimit -n
65536
```

Notes

- If the correct number of open files has not been established, reboot the host:


```
# sync
# reboot
```

 If the settings still do not take effect, enable UsePAM (set value to yes) in `/etc/ssh/sshd_config` and restart SSH again.
- If the `nofile` parameter is less than 32768, HP Vertica fails and logs the following message:


```
Not enough open file handles allowed (1024 available/32768 required);
see 'ulimit -n'.
```

See also

Verify that kernel and user parameters were set (page [63](#))

Linux man proc page <http://linux.die.net/man/5/proc>

Increase max_map_count Parameter

Note: This step is automatically done by the installer. See **Verify that Kernel and User Parameters Were Set** (page [63](#)) for details on the setting. You only need to perform this step if the automatic setting did not work.

Increase the `max_map_count` kernel parameter to avoid running out of mapped areas, especially on nodes with lots of memory.

- 1 Determine the amount of total memory in KB:

```
$ grep MemTotal /proc/meminfo | awk '{printf "%.0f", $2}'
```

- 2 Take that number and divide by 16. That is the *VALUE* to use below.

Note: The default value is 65536, which might be too low for HP Vertica to run. At a minimum, set `max_map_count` to the value you determined above, or 65536, whichever is higher.

- 3 Issue the following command as root:

```
$ sysctl vm.max_map_count=VALUE
```

The above changes will be reverted on machine reboot. To make the change permanent, as root, issue:

```
echo "vm.max_map_count = VALUE" | tee -a /etc/sysctl.conf
```

- 4 Reload the config file as root:

```
$ sysctl -p
```

- 5 Log out and log in again.

Set Up Time Synchronization

Use NTP or some other time-synchronization mechanism to keep all hosts synchronized. Time variances can cause inconsistent query results when using Date/Time Functions or when using a network authentication protocol, like Kerberos.

To verify that the Network Time Protocol Daemon (NTPD) is operating correctly, issue the following command on all nodes in the cluster:

```
/usr/sbin/ntpq -c rv | grep stratum
```

A stratum level of 16 indicates that NTP is not synchronizing correctly.

See also

- To configure the system time, date, and the Network Time Protocol in the graphical interface:
<http://kbase.redhat.com/faq/docs/DOC-6731>
<http://kbase.redhat.com/faq/docs/DOC-6731>
- To synchronize the system clock to Network Time Protocol:
<http://kbase.redhat.com/faq/docs/DOC-6902>
<http://kbase.redhat.com/faq/docs/DOC-6902>
- To synchronize the system clock to a Network Time Protocol server immediately:
<http://kbase.redhat.com/faq/docs/DOC-6991>
<http://kbase.redhat.com/faq/docs/DOC-6991>

Note: The above links were current as of the last publication of the HP Vertica documentation and could change between releases.

Manage Large Grouped ROS Containers

If you upgrade your database from HP Vertica 4.0.x to 4.0 Patchset 1, check for the existence of large grouped ROS containers, especially if you used COPY with the AUTO option, or without specifying any load option. AUTO is the default option for loading bulk data with COPY, as described in the COPY statement in the SQL Reference Manual. Note that this issue does not affect a 3.5.x database.

Resolution

Small unpartitioned tables can be merged entirely using the following statement:

```
=> SELECT DO_TM_TASK('mergeout', 'table-name');
```

On large partitioned tables, re-merge partitions with grouped ROS containers using the MERGE_PARTITIONS() command. This command eliminates grouped ROSes.

1 Create two views to identify grouped ROS containers that require re-merging:

```
=> CREATE VIEW large_grouped_roses AS SELECT * FROM (
  SELECT s.*, TO_CHAR(100.0*grouped_bytes::FLOAT/
    (grouped_bytes::FLOAT+NON_GROUPED_BYTES), '999.99') AS PCT
  FROM (
    SELECT sc.schema_name, sc.projection_name,
      SUM(CASE WHEN GROUPING = 'ALL' THEN USED_BYTES ELSE 0 END)
        AS GROUPED_BYTES,
      SUM(CASE WHEN GROUPING = 'ALL' THEN 0 ELSE USED_BYTES END)
        AS NON_GROUPED_BYTES
    FROM storage_containers SC
    GROUP BY sc.schema_name, sc.projection_name) AS S ) AS Q
  WHERE q.grouped_bytes > 1024000 AND Q.PCT > 10
  ORDER BY Q.GROUPED_BYTES DESC;
=> CREATE VIEW proj_to_merge AS SELECT distinct schema_name,
  projection_name, partition_key
  FROM (
    SELECT gr.schema_name, gr.projection_name, sc.storage_oid,
      used_bytes, partition_key
    FROM large_grouped_roses GR
    NATURAL LEFT JOIN storage_containers SC
    LEFT JOIN partitions P
    ON sc.storage_oid = p.ros_id
    WHERE sc.grouping = 'ALL') Q;
```

2 Execute the following command to identify any large grouped ROS containers:

```
=> SELECT * FROM large_grouped_roses;
```

The following example output identifies two projections:

```
-[ RECORD 1 ]-----+-----
schema_name      | public
projection_name   | product_dimension_DBD_14_seg_vmartdb_design
grouped_bytes     | 1979034
```

```
non_grouped_bytes | 0
pct                | 100.00
-[ RECORD 2 ]-----+-----
schema_name       | public
projection_name    | product_dimension_DBD_2_seg_vmartdb_design
grouped_bytes      | 1979034
non_grouped_bytes  | 0
pct                | 100.00
```

- 3** Execute the following command to identify the partition key that belongs to the containers you want to merge:

```
=> SELECT * FROM proj_to_merge;
```

If the table is partitioned, the results should return the affected partition key. If the table is not partitioned, as in the following example, no partition key is returned:

```
-[ RECORD 1 ]---+-----
schema_name     | public
projection_name | product_dimension_DBD_14_seg_vmartdb_design
partition_key    |
-[ RECORD 2 ]---+-----
schema_name     | public
projection_name | product_dimension_DBD_2_seg_vmartdb_design
partition_key    |
```

- 4** Do one of the following:

1. If a partition_key is listed, execute the following command on the projection/partition:

```
=> SELECT merge_partitions('table_name', 'from_key', 'to_key');
```

2. If no partition key is returned, execute the following command:

```
=> SELECT DO_TM_TASK('mergeout');
DO_TM_TASK | mergeout for projection 'product_dimension_node0001'
mergeout for projection 'product_dimension_node0002'
mergeout for projection 'product_dimension_node0003'
mergeout for projection 'product_dimension_node0004'
mergeout for projection
'call_center_dimension_DBD_32_seg_vmartdb_design'
...
```

IMPORTANT: Do not perform step 4b on large partitioned tables.

See also

DO_TM_TASK and MERGE_PARTITIONS in the SQL Reference Manual

Appendix: Time Zones

Using Time Zones With HP Vertica

HP Vertica uses the `TZ` environment variable on each node, if it has been set, for the default current time zone. Otherwise, HP Vertica uses the operating system time zone.

The `TZ` variable can be set by the operating system during login (see `/etc/profile`, `/etc/profile.d`, or `/etc/bashrc`) or by the user in `.profile`, `.bashrc` or `.bash-profile`.

`TZ` must be set to the same value on each node when you start HP Vertica.

The following command returns the current time zone for your database:

```
=> SHOW TIMEZONE;
  name  |      setting
-----+-----
 timezone | America/New_York
(1 row)
```

You can also use the `SET TIMEZONE TO { value | 'value' }` command to set the time zone for a single session.

There is no database default time zone; instead, `TIMESTAMP WITH TIMEZONE` (`TIMESTAMPTZ`) data is stored in GMT (UTC) by converting data from the current local time zone to GMT.

When `TIMESTAMPTZ` data is used, data is converted back to use the current local time zone, which might be different from the local time zone where the data was stored. This conversion takes into account Daylight Saving Time (Summer Time), if applicable, depending on the year and date, to know when the Daylight Saving Time change occurred.

`TIMESTAMP WITHOUT TIMEZONE` data stores the timestamp, as given, and retrieves it exactly as given. The current time zone is ignored. The same is true for `TIME WITHOUT TIMEZONE`. For `TIME WITH TIMEZONE` (`TIMETZ`), however, the current time zone setting is stored along with the given time, and that time zone is used on retrieval.

Note: HP recommends that you use `TIMESTAMPTZ`, not `TIMETZ`.

`TIMESTAMPTZ` uses the current time zone on both input and output, such as in the following example:

```
=> CREATE TEMP TABLE s (tstz TIMESTAMPTZ);
=> SET TIMEZONE TO 'America/New_York';
=> INSERT INTO s VALUES ('2009-02-01 00:00:00');
=> INSERT INTO s VALUES ('2009-05-12 12:00:00');
=> SELECT tstz AS 'Local timezone', tstz AT TIMEZONE 'America/New_York' AS
   'America/New_York',
   tstz AT TIMEZONE 'GMT' AS 'GMT' FROM s;
```

Local timezone	America/New_York	GMT
2009-02-01 00:00:00-05	2009-02-01 00:00:00	2009-02-01 05:00:00
2009-05-12 12:00:00-04	2009-05-12 12:00:00	2009-05-12 16:00:00

(2 rows)

The -05 in the Local time zone column above shows that the data is displayed in EST, while -04 indicates EDT. The other two columns show the `TIMESTAMP WITHOUT TIMEZONE` at the specified time zone.

The next example illustrates what occurs if the current time zone is changed to, for example, Greenwich Mean Time:

```
=> SET TIMEZONE TO 'GMT';
=> SELECT tstz AS 'Local timezone', tstz AT TIMEZONE 'America/New_York' AS
    'America/New_York', tstz AT TIMEZONE 'GMT' as 'GMT' FROM s;
```

Local timezone	America/New_York	GMT
2009-02-01 05:00:00+00	2009-02-01 00:00:00	2009-02-01 05:00:00
2009-05-12 16:00:00+00	2009-05-12 12:00:00	2009-05-12 16:00:00

(2 rows)

The +00 in the Local time zone column above indicates that `TIMESTAMPTZ` is displayed in 'GMT'.

The approach of using `TIMESTAMPTZ` fields to record events captures the GMT of the event, as expressed in terms of the local time zone. Later, it allows for easy conversion to any other time zone, either by setting the local time zone or by specifying an explicit `AT TIMEZONE` clause.

The following example shows how `TIMESTAMP WITHOUT TIMEZONE` fields work in HP Vertica.

```
=> CREATE TEMP TABLE tnoz (ts TIMESTAMP);
=> INSERT INTO tnoz VALUES('2009-02-01 00:00:00');
=> INSERT INTO tnoz VALUES('2009-05-12 12:00:00');
=> SET TIMEZONE TO 'GMT';
=> SELECT ts AS 'No timezone', ts AT TIMEZONE 'America/New_York' AS
    'America/New_York', ts AT TIMEZONE 'GMT' AS 'GMT' FROM tnoz;
```

No timezone	America/New_York	GMT
2009-02-01 00:00:00	2009-02-01 05:00:00+00	2009-02-01 00:00:00+00
2009-05-12 12:00:00	2009-05-12 16:00:00+00	2009-05-12 12:00:00+00

(2 rows)

The +00 at the end of a timestamp indicates that the setting is `TIMESTAMP WITH TIMEZONE` in GMT (the current time zone). The 'America/New_York' column shows what the 'GMT' setting was when you recorded the time, assuming you read a normal clock in the time zone 'America/New_York'. What this shows is that if it is midnight in the 'America/New_York' time zone, then it is 5 am GMT.

Note: 00:00:00 Sunday February 1, 2009 in America/New_York converts to 05:00:00 Sunday February 1, 2009 in GMT.

The 'GMT' column displays the GMT time, assuming the input data was captured in GMT.

If you don't set the time zone to GMT, and you use another time zone, for example 'America/New_York', then the results display in 'America/New_York' with a -05 and -04, showing the difference between that time zone and GMT.

```
=> SET TIMEZONE TO 'America/New_York';
=> SHOW TIMEZONE;
      name      |      setting
-----+-----
      timezone | America/New_York
(1 row)

=> SELECT ts AS 'No timezone', ts AT TIMEZONE 'America/New_York' AS
      'America/New_York', ts AT TIMEZONE 'GMT' AS 'GMT' FROM tnoz;
      No timezone      |      America/New_York      |      GMT
-----+-----+-----
2009-02-01 00:00:00 | 2009-02-01 00:00:00-05 | 2009-01-31 19:00:00-05
2009-05-12 12:00:00 | 2009-05-12 12:00:00-04 | 2009-05-12 08:00:00-04
(2 rows)
```

In this case, the last column is interesting in that it returns the time in New York, given that the data was captured in 'GMT'.

See also

Set the default time zone (page [21](#))

SET TIMEZONE and Date/Time Data Types in the SQL Reference Manual

Africa

Africa/Abidjan
Africa/Accra
Africa/Addis_Ababa
Africa/Algiers
Africa/Asmera
Africa/Bamako
Africa/Bangui
Africa/Banjul
Africa/Bissau
Africa/Blantyre

Africa/Brazzaville
Africa/Bujumbura
Africa/Cairo Egypt
Africa/Casablanca
Africa/Ceuta
Africa/Conakry
Africa/Dakar
Africa/Dar_es_Salaam
Africa/Djibouti
Africa/Douala
Africa/El_Aaiun
Africa/Freetown
Africa/Gaborone
Africa/Harare
Africa/Johannesburg
Africa/Kampala
Africa/Khartoum
Africa/Kigali
Africa/Kinshasa
Africa/Lagos
Africa/Libreville
Africa/Lome
Africa/Luanda
Africa/Lubumbashi

Africa/Lusaka
Africa/Malabo
Africa/Maputo
Africa/Maseru
Africa/Mbabane
Africa/Mogadishu
Africa/Monrovia
Africa/Nairobi
Africa/Ndjamena
Africa/Niamey
Africa/Nouakchott
Africa/Ouagadougou
Africa/Porto-Novo
Africa/Sao_Tome
Africa/Timbuktu
Africa/Tripoli Libya
Africa/Tunis
Africa/Windhoek

America

America/Adak America/Atka US/Aleutian
America/Anchorage SystemV/YST9YDT US/Alaska
America/Anguilla

America/Antigua
America/Araguaina
America/Aruba
America/Asuncion
America/Bahia
America/Barbados
America/Belem
America/Belize
America/Boa_Vista
America/Bogota
America/Boise
America/Buenos_Aires
America/Cambridge_Bay
America/Campo_Grande
America/Cancun
America/Caracas
America/Catamarca
America/Cayenne
America/Cayman
America/Chicago CST6CDT SystemV/CST6CDT US/Central
America/Chihuahua
America/Cordoba America/Rosario
America/Costa_Rica

America/Cuiaba
America/Curacao
America/Danmarkshavn
America/Dawson
America/Dawson_Creek
America/Denver MST7MDT SystemV/MST7MDT US/Mountain America/Shiprock Navajo
America/Detroit US/Michigan
America/Dominica
America/Edmonton Canada/Mountain
America/Eirunepe
America/El_Salvador
America/Ensenada America/Tijuana Mexico/BajaNorte
America/Fortaleza
America/Glace_Bay
America/Godthab
America/Goose_Bay
America/Grand_Turk
America/Grenada
America/Guadeloupe
America/Guatemala
America/Guayaquil
America/Guyana
America/Halifax Canada/Atlantic SystemV/AST4ADT

America/Havana Cuba
America/Hermosillo
America/Indiana/Indianapolis America/Indianapolis America/Fort_Wayne EST SystemV/EST5 US/East-Indiana
America/Indiana/Knox America/Knox_IN US/Indiana-Starke
America/Indiana/Marengo
America/Indiana/Vevay
America/Inuvik
America/Iqaluit
America/Jamaica Jamaica
America/Jujuy
America/Juneau
America/Kentucky/Louisville America/Louisville
America/Kentucky/Monticello
America/La_Paz
America/Lima
America/Los_Angeles PST8PDT SystemV/PST8PDT US/Pacific US/Pacific- New
America/Maceio
America/Managua
America/Manaus Brazil/West
America/Martinique
America/Mazatlan Mexico/BajaSur

America/Mendoza
America/Menominee
America/Merida
America/Mexico_City Mexico/General
America/Miquelon
America/Monterrey
America/Montevideo
America/Montreal
America/Montserrat
America/Nassau
America/New_York EST5EDT SystemV/EST5EDT US/Eastern
America/Nipigon
America/Nome
America/Noronha Brazil/DeNoronha
America/North_Dakota/Center
America/Panama
America/Pangnirtung
America/Paramaribo
America/Phoenix MST SystemV/MST7 US/Arizona
America/Port-au-Prince
America/Port_of_Spain
America/Porto_Acre America/Rio_Branco Brazil/Acre
America/Porto_Velho

America/Puerto_Rico SystemV/AST4
America/Rainy_River
America/Rankin_Inlet
America/Recife
America/Regina Canada/East-Saskatchewan Canada/Saskatchewan SystemV/CST6
America/Santiago Chile/Continental
America/Santo_Domingo
America/Sao_Paulo Brazil/East
America/Scoresbysund
America/St_Johns Canada/Newfoundland
America/St_Kitts
America/St_Lucia
America/St_Thomas America/Virgin
America/St_Vincent
America/Swift_Current
America/Tegucigalpa
America/Thule
America/Thunder_Bay
America/Toronto Canada/Eastern
America/Tortola
America/Vancouver Canada/Pacific
America/Whitehorse Canada/Yukon
America/Winnipeg Canada/Central

America/Yakutat
America/Yellowknife

Antarctica

Antarctica/Casey
Antarctica/Davis
Antarctica/DumontDUrville
Antarctica/Mawson
Antarctica/McMurdo
Antarctica/South_Pole
Antarctica/Palmer
Antarctica/Rothera
Antarctica/Syowa
Antarctica/Vostok

Asia

Asia/Aden
Asia/Almaty
Asia/Amman
Asia/Anadyr
Asia/Aqtau
Asia/Aqtobe

Asia/Ashgabat Asia/Ashkhabad
Asia/Baghdad
Asia/Bahrain
Asia/Baku
Asia/Bangkok
Asia/Beirut
Asia/Bishkek
Asia/Brunei
Asia/Calcutta
Asia/Choibalsan
Asia/Chongqing Asia/Chungking
Asia/Colombo
Asia/Dacca Asia/Dhaka
Asia/Damascus
Asia/Dili
Asia/Dubai
Asia/Dushanbe
Asia/Gaza
Asia/Harbin
Asia/Hong_Kong Hongkong
Asia/Hovd
Asia/Irkutsk
Asia/Jakarta
Asia/Jayapura

Asia/Jerusalem Asia/Tel_Aviv Israel
Asia/Kabul
Asia/Kamchatka
Asia/Karachi
Asia/Kashgar
Asia/Katmandu
Asia/Krasnoyarsk
Asia/Kuala_Lumpur
Asia/Kuching
Asia/Kuwait
Asia/Macao Asia/Macau
Asia/Magadan
Asia/Makassar Asia/Ujung_Pandang
Asia/Manila
Asia/Muscat
Asia/Nicosia Europe/Nicosia
Asia/Novosibirsk
Asia/Omsk
Asia/Oral
Asia/Phnom_Penh
Asia/Pontianak
Asia/Pyongyang
Asia/Qatar
Asia/Qyzylorda

Asia/Rangoon
Asia/Riyadh
Asia/Riyadh87 Mideast/Riyadh87
Asia/Riyadh88 Mideast/Riyadh88
Asia/Riyadh89 Mideast/Riyadh89
Asia/Saigon
Asia/Sakhalin
Asia/Samarkand
Asia/Seoul ROK
Asia/Shanghai PRC
Asia/Singapore Singapore
Asia/Taipei ROC
Asia/Tashkent
Asia/Tbilisi
Asia/Tehran Iran
Asia/Thimbu Asia/Thimphu
Asia/Tokyo Japan
Asia/Ulaanbaatar Asia/Ulan_Bator
Asia/Urumqi
Asia/Vientiane
Asia/Vladivostok
Asia/Yakutsk
Asia/Yekaterinburg
Asia/Yerevan

Atlantic

Atlantic/Azores
Atlantic/Bermuda
Atlantic/Canary
Atlantic/Cape_Verde
Atlantic/Faeroe
Atlantic/Madeira
Atlantic/Reykjavik Iceland
Atlantic/South_Georgia
Atlantic/St_Helena
Atlantic/Stanley

Australia

Australia/ACT
Australia/Canberra
Australia/NSW
Australia/Sydney
Australia/Adelaide
Australia/South
Australia/Brisbane
Australia/Queensland
Australia/Broken_Hill

Australia/Yancowinna
Australia/Darwin Australia/North
Australia/Hobart Australia/Tasmania
Australia/LHI Australia/Lord_Howe
Australia/Lindeman
Australia/Melbourne Australia/Victoria
Australia/Perth Australia/West

Etc/GMT

Etc/GMT+1
Etc/GMT+2
Etc/GMT+3
Etc/GMT+4
Etc/GMT+5
Etc/GMT+6
Etc/GMT+7
Etc/GMT+8
Etc/GMT+9
Etc/GMT+10
Etc/GMT+11

Etc/GMT+12
Etc/GMT-1
Etc/GMT-2
Etc/GMT-3
Etc/GMT-4
Etc/GMT-5
Etc/GMT-6
Etc/GMT-7
Etc/GMT-8
Etc/GMT-9
Etc/GMT-10
Etc/GMT-11
Etc/GMT-12
Etc/GMT-13
Etc/GMT-14

Europe

Europe/Amsterdam
Europe/Andorra
Europe/Athens
Europe/Belfast
Europe/Belgrade
Europe/Ljubljana

Europe/Sarajevo
Europe/Skopje
Europe/Zagreb
Europe/Berlin
Europe/Brussels
Europe/Bucharest
Europe/Budapest
Europe/Chisinau Europe/Tiraspol
Europe/Copenhagen
Europe/Dublin Eire
Europe/Gibraltar
Europe/Helsinki
Europe/Istanbul Asia/Istanbul Turkey
Europe/Kaliningrad
Europe/Kiev
Europe/Lisbon Portugal
Europe/London GB GB-Eire
Europe/Luxembourg
Europe/Madrid
Europe/Malta
Europe/Minsk
Europe/Monaco
Europe/Moscow W-SU
Europe/Oslo
Arctic/Longyearbyen

Atlantic/Jan_Mayen
Europe/Paris
Europe/Prague Europe/Bratislava
Europe/Riga
Europe/Rome Europe/San_Marino Europe/Vatican
Europe/Samara
Europe/Simferopol
Europe/Sofia
Europe/Stockholm
Europe/Tallinn
Europe/Tirane
Europe/Uzhgorod
Europe/Vaduz
Europe/Vienna
Europe/Vilnius
Europe/Warsaw Poland
Europe/Zaporozhye
Europe/Zurich

Indian

Indian/Antananarivo
Indian/Chagos
Indian/Christmas

Indian/Cocos
Indian/Comoro
Indian/Kerguelen
Indian/Mahe
Indian/Maldives
Indian/Mauritius
Indian/Mayotte
Indian/Reunion

Pacific

Pacific/Apia
Pacific/Auckland NZ
Pacific/Chatham NZ-CHAT
Pacific/Easter Chile/EasterIsland
Pacific/Efate
Pacific/Enderbury
Pacific/Fakaofu
Pacific/Fiji
Pacific/Funafuti
Pacific/Galapagos
Pacific/Gambier System V/YST9
Pacific/Guadalcanal

Pacific/Guam
Pacific/Honolulu HST SystemV/HST10 US/Hawaii
Pacific/Johnston
Pacific/Kiritimati
Pacific/Kosrae
Pacific/Kwajalein Kwajalein
Pacific/Majuro
Pacific/Marquesas
Pacific/Midway
Pacific/Nauru
Pacific/Niue
Pacific/Norfolk
Pacific/Noumea
Pacific/Pago_Pago
Pacific/Samoa US/Samoa
Pacific/Palau
Pacific/Pitcairn SystemV/PST8
Pacific/Ponape
Pacific/Port_Moresby
Pacific/Rarotonga
Pacific/Saipan
Pacific/Tahiti
Pacific/Tarawa
Pacific/Tongatapu

Pacific/Truk
Pacific/Wake
Pacific/Wallis
Pacific/Yap

Copyright Notice

Copyright© 2006-2013 Hewlett-Packard, and its licensors. All rights reserved.

Hewlett-Packard 150 CambridgePark Drive Cambridge, MA 02140 Phone: +1 617 386 4400 E-Mail: info@vertica.com Web site: http://www.vertica.com (http://www.vertica.com)

The software described in this copyright notice is furnished under a license and may be used or copied only in accordance with the terms of such license. Hewlett-Packard software contains proprietary information, as well as trade secrets of Hewlett-Packard, and is protected under international copyright law. Reproduction, adaptation, or translation, in whole or in part, by any means — graphic, electronic or mechanical, including photocopying, recording, taping, or storage in an information retrieval system — of any part of this work covered by copyright is prohibited without prior written permission of the copyright owner, except as allowed under the copyright laws.

This product or products depicted herein may be protected by one or more U.S. or international patents or pending patents.

Trademarks

HP Vertica™, the HP Vertica Analytics Platform™, and FlexStore™ are trademarks of Hewlett-Packard.

Adobe®, Acrobat®, and Acrobat® Reader® are registered trademarks of Adobe Systems Incorporated.

AMD™ is a trademark of Advanced Micro Devices, Inc., in the United States and other countries.

DataDirect® and DataDirect Connect® are registered trademarks of Progress Software Corporation in the U.S. and other countries.

Fedora™ is a trademark of Red Hat, Inc.

Intel® is a registered trademark of Intel.

Linux® is a registered trademark of Linus Torvalds.

Microsoft® is a registered trademark of Microsoft Corporation.

Novell® is a registered trademark and SUSE™ is a trademark of Novell, Inc., in the United States and other countries.

Oracle® is a registered trademark of Oracle Corporation.

Red Hat® is a registered trademark of Red Hat, Inc.

VMware® is a registered trademark or trademark of VMware, Inc., in the United States and/or other jurisdictions.

Other products mentioned may be trademarks or registered trademarks of their respective companies.

Information on third-party software used in HP Vertica, including details on open-source software, is available in the guide [Third-Party Software Acknowledgements](#).