

---

Vertica Database 2.0.5-0

# Troubleshooting Guide

Copyright© 2006, 2007 Vertica Systems, Inc.

Date of Publication: 1/8/2008



CONFIDENTIAL

# Copyright Notice

---

Copyright© 2006 - 2007 Vertica Systems, Inc. and its licensors. All rights reserved.

Vertica Systems, Inc. Three Dundee Park Drive, Suite 102 Andover, MA 01810-3723 Phone: (978) 475-1070 Fax: (978) 475-6855 E-Mail: <a href="mailto:info@vertica.com">info@vertica.com</a> Web site: <a href="http://www.vertica.com">http://www.vertica.com</a> ( <a href="http://www.vertica.com">http://www.vertica.com</a> )
---

The software described in this document is furnished under a license and may be used or copied only in accordance with the terms of such license. Vertica Systems, Inc. software contains proprietary information, as well as trade secrets of Vertica Systems, Inc., 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

Vertica™ and the Vertica Database™ are trademarks of Vertica Systems, Inc.

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.

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.

## Open Source Software Acknowledgements

## Boost

Boost Software License - Version 1.0 - August 17th, 2003

Permission is hereby granted, free of charge, to any person or organization obtaining a copy of the software and accompanying documentation covered by this license (the "Software") to use, reproduce, display, distribute, execute, and transmit the Software, and to prepare derivative works of the Software, and to permit third-parties to whom the Software is furnished to do so, all subject to the following:

The copyright notices in the Software and this entire statement, including the above license grant, this restriction and the following disclaimer, must be included in all copies of the Software, in whole or in part, and all derivative works of the Software, unless such copies or derivative works are solely in the form of machine-executable object code generated by a source language processor.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR ANYONE DISTRIBUTING THE SOFTWARE BE LIABLE FOR ANY DAMAGES OR OTHER LIABILITY, WHETHER IN CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

## PostgreSQL

This product uses the PostgreSQL Database Management System(formerly known as Postgres, then as Postgres95)

Portions Copyright © 1996-2005, The PostgreSQL Global Development Group

Portions Copyright © 1994, The Regents of the University of California

Permission to use, copy, modify, and distribute this software and its documentation for any purpose, without fee, and without a written agreement is hereby granted, provided that the above copyright notice and this paragraph and the following two paragraphs appear in all copies.

IN NO EVENT SHALL THE UNIVERSITY OF CALIFORNIA BE LIABLE TO ANY PARTY FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, INCLUDING LOST PROFITS, ARISING OUT OF THE USE OF THIS SOFTWARE AND ITS DOCUMENTATION, EVEN IF THE UNIVERSITY OF CALIFORNIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

THE UNIVERSITY OF CALIFORNIA SPECIFICALLY DISCLAIMS ANY WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE SOFTWARE PROVIDED HEREUNDER IS ON AN "AS IS" BASIS, AND THE UNIVERSITY OF CALIFORNIA HAS NO OBLIGATIONS TO PROVIDE MAINTENANCE, SUPPORT, UPDATES, ENHANCEMENTS, OR MODIFICATIONS.

## Python Dialog

The Administration Tools part of this product uses Python Dialog,a Python module for doing console-mode user interaction.

Upstream Author:

Peter Astrand <[peter@cendio.se](mailto:peter@cendio.se)>

Robb Shecter <[robb@acm.org](mailto:robb@acm.org)>

Sultanbek Tezadov <<http://sultan.da.ru>>

Florent Rougon <[flo@via.ecp.fr](mailto:flo@via.ecp.fr)>

Copyright © 2000 Robb Shecter, Sultanbek Tezadov

Copyright © 2002, 2003, 2004 Florent Rougon

License:

This package is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This package is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU Lesser General Public License for more details.

You should have received a copy of the GNU Lesser General Public License along with this package; if not, write to the Free Software Foundation, Inc., 51 Franklin St, Fifth Floor, Boston, MA 02110-1301 USA

On Vertica systems, complete source code of the Python dialog package and complete text of the GNU Lesser General Public License can be found on the Vertica Systems website at <http://www.vertica.com/licenses/pythondialog-2.7.tar.bz2> <http://www.vertica.com/licenses/pythondialog-2.7.tar.bz2>

## Spread

This product uses software developed by Spread Concepts LLC for use in the Spread toolkit. For more information about Spread see <http://www.spread.org> (<http://www.spread.org>).

Copyright (c) 1993-2006 Spread Concepts LLC. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer and request.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer and request in the documentation and/or other materials provided with the distribution.
3. All advertising materials (including web pages) mentioning features or use of this software, or software that uses this software, must display the following acknowledgment: "This product uses software developed by Spread Concepts LLC for use in the Spread toolkit. For more information about Spread see <http://www.spread.org>"
4. The names "Spread" or "Spread toolkit" must not be used to endorse or promote products derived from this software without prior written permission.
5. Redistributions of any form whatsoever must retain the following acknowledgment:

"This product uses software developed by Spread Concepts LLC for use in the Spread toolkit. For more information about Spread, see <http://www.spread.org>"

6. This license shall be governed by and construed and enforced in accordance with the laws of the State of Maryland, without reference to its conflicts of law provisions. The exclusive jurisdiction and venue for all legal actions relating to this license shall be in courts of competent subject matter jurisdiction located in the State of Maryland.

TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, SPREAD IS PROVIDED UNDER THIS LICENSE ON AN AS IS BASIS, WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, WARRANTIES THAT SPREAD IS FREE OF DEFECTS,

MERCHANTABLE, FIT FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT. ALL WARRANTIES ARE DISCLAIMED AND THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE CODE IS WITH YOU. SHOULD ANY CODE PROVE DEFECTIVE IN ANY RESPECT, YOU (NOT THE COPYRIGHT HOLDER OR ANY OTHER CONTRIBUTOR) ASSUME THE COST OF ANY NECESSARY SERVICING, REPAIR OR CORRECTION. THIS DISCLAIMER OF WARRANTY CONSTITUTES AN ESSENTIAL PART OF THIS LICENSE. NO USE OF ANY CODE IS AUTHORIZED HEREUNDER EXCEPT UNDER THIS DISCLAIMER.

TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, IN NO EVENT SHALL THE COPYRIGHT HOLDER OR ANY OTHER CONTRIBUTOR BE LIABLE FOR ANY SPECIAL, INCIDENTAL, INDIRECT, OR CONSEQUENTIAL DAMAGES FOR LOSS OF PROFITS, REVENUE, OR FOR LOSS OF INFORMATION OR ANY OTHER LOSS.

YOU EXPRESSLY AGREE TO FOREVER INDEMNIFY, DEFEND AND HOLD HARMLESS THE COPYRIGHT HOLDERS AND CONTRIBUTORS OF SPREAD AGAINST ALL CLAIMS, DEMANDS, SUITS OR OTHER ACTIONS ARISING DIRECTLY OR INDIRECTLY FROM YOUR ACCEPTANCE AND USE OF SPREAD.

Although NOT REQUIRED, we at Spread Concepts would appreciate it if active users of Spread put a link on their web site to Spread's web site when possible. We also encourage users to let us know who they are, how they are using Spread, and any comments they have through either e-mail ([spread@spread.org](mailto:spread@spread.org)) or our web site at (<http://www.spread.org/comments>).



# Contents

<b>Copyright Notice</b>	<b>ii</b>
-------------------------	-----------

---

<b>Technical Support</b>	<b>11</b>
--------------------------	-----------

---

<b>About the Documentation</b>	<b>13</b>
--------------------------------	-----------

---

Where to Find the Vertica Documentation .....	13
Reading the HTML Files .....	14
Printing the PDF Files.....	14

<b>Suggested Reading Paths</b>	<b>15</b>
--------------------------------	-----------

---

Where to Find Additional Information .....	17
Typographical Conventions .....	18

<b>Preface</b>	<b>19</b>
----------------	-----------

---

<b>Overview</b>	<b>21</b>
-----------------	-----------

---

<b>Monitoring the Database</b>	<b>22</b>
--------------------------------	-----------

---

View Database Cluster State.....	22
Monitoring the Log Files .....	23
Rotating the Log Files.....	24
Using the SQL Monitoring API.....	25
Monitoring Processes.....	26
Monitoring System Resource Usage.....	26

<b>Using Diagnostic Tools</b>	<b>29</b>
-------------------------------	-----------

---

Determining the Version Number.....	30
Diagnostics Utility .....	30
DUMP_CATALOG.....	30
DUMP_LOCKTABLE .....	31

<b>Failure Recovery</b>	<b>31</b>
-------------------------	-----------

---

Shutdown Problems .....	32
Cannot shut down while users are connected .....	32
Database ... did not appear to stop.....	33

## Troubleshooting Guide

---

Startup Problems.....	33
Startup successful, but some nodes are recovering.....	34
Error starting database, no nodes are up.....	34
Database startup successful, but it may be incomplete.....	35
Database did not start cleanly on initiator node!.....	35
TIMEOUT ERROR: Could not login with SSH.....	36
Good epoch logs are available on all nodes.....	36
No good epoch log available on node.....	37
Nodes stuck in INITIALIZING state.....	37
Spread is not running.....	38

## Installation Problems 39

---

Failed Dependencies.....	39
SSH already configured.....	40

## Configuration Problems 41

---

Administration Tools Errors.....	41
Cannot find IP address for host.....	41
Clock skew detected.....	41
Could not login with SSH.....	42
Could not open configuration file.....	42
Create Database Fails.....	43
Database Designer Errors.....	45
No such file or directory ... super_projections.tmp.sql.....	45



---

Error detected during node verification .....	45
Not enough open file handles allowed.....	45
Permission denied .....	46
Spread Panic .....	46
<b>Connection Problems</b>	<b>47</b>

---

could not connect to database .....	47
could not connect to server: connection refused .....	47
Firewall Configuration.....	47
<b>SQL Execution Problems</b>	<b>49</b>

---

Column Definition Changed to NOT NULL .....	49
CREATE PROJECTION Hangs.....	49
Ctrl-C Distorts Query Result Set Output .....	50
DEPRECATED syntax .....	50
INSERT Fails.....	50
Insufficient Projections to answer query.....	50
LIKE Predicate Referred to as ~~ .....	51
Missing FROM-clause entry for table .....	51
Not a Star or Snowflake Query .....	52
Out of Memory .....	52
Query Performance .....	53
Site <node-name> got error: Too many groups. ....	53
Unexpected Query Results.....	54
Update and Delete Performance .....	54
WOS Overload.....	54
<b>Platform Problems</b>	<b>55</b>

---

Data Disk Spontaneously Remounts.....	55
<b>Index</b>	<b>57</b>

---



# Technical Support

---

To submit problem reports, questions, comments, and suggestions, please use the Technical Support page on the Vertica Systems, Inc. Web site:

<http://www.vertica.com/support> (http://www.vertica.com/support)

You must be a registered user in order to access the page.

Before reporting a problem, please run the **Diagnostics Utility** (page 30) described in the Troubleshooting Guide and attach the resulting .zip file.



# About the Documentation

---

## Where to Find the Vertica Documentation

Vertica Systems, Inc. recommends that you copy the Vertica documentation from the database server (any cluster host) to a client system on which you can use a browser and/or the Adobe Reader.

### Database Server Systems

The Vertica Database Documentation Set is automatically installed in the `/opt/vertica/doc/` directory on all cluster hosts. If you have a browser and/or the Adobe Reader installed on a cluster host, you can access the documentation directly.

<code>/opt/vertica/doc/HTML/Master/index.htm</code>
<code>/opt/vertica/doc/JDBC/index.htm</code>
<code>/opt/vertica/doc/PDF/book-name.pdf</code>

### Database Client Systems

To create a copy of the Vertica documentation on a client system, do one of the following:

- Download the documentation package (`.tar.gz` or `.zip`) from the Vertica Systems, Inc. Web site and extract the files to a directory on the client system, using the original pathnames
- Copy the documentation directories from a database server system to a convenient location in your client system. All cross-references within the HTML documentation are relative so there is no location dependency.

<code>&lt;your-location&gt;HTML/Master/index.htm</code>
<code>&lt;your-location&gt;JDBC/index.htm</code>
<code>&lt;your-location&gt;PDF/book-name.htm</code>

### World Wide Web

You can read and/or download the Vertica documentation from the [Vertica Systems, Inc. Web site](http://www.vertica.com/v-zone/product_documentation):

[http://www.vertica.com/v-zone/product\\_documentation](http://www.vertica.com/v-zone/product_documentation) `http://www.vertica.com/v-zone/product_documentation`.

You need a V-Zone login to access the Product Documentation page.

The documentation on the Vertica Systems, Inc. Web site is updated each time a new release is issued. If you are using an older version, refer to the documentation on your database server or client systems.

## Reading the HTML Files

The Vertica documentation files are provided in HTML browser format for platform independence. The HTML files only require a browser that can **display frames** properly and has **JavaScript** enabled. The HTML files **do not require a Web (HTTP) server**.

The Vertica documentation has been tested on the following browsers:

- Internet Explorer 7
- FireFox
- Opera
- Safari

Please report any script, image rendering, or text formatting problems to **Technical Support** (on page 11).

The Vertica documentation may contain links to Web sites of other companies or organizations that Vertica does not own or control. If any of these links are broken, please inform us.

## Printing the PDF Files

The documentation files are supplied in **Adobe Acrobat™ PDF** document format for the purpose of making printed copies as needed. The documents are designed to be printed on standard 8½ x 11 paper using full duplex (two sided printing).

You can open and print any of the PDF documents using the **Adobe Reader**. (You can download the latest version of the free Acrobat Reader from the **Adobe Web site** (<http://www.adobe.com/products/acrobat/readstep2.html>).

HTML links to the PDF files are provided here for browser access.

- Database Administrator's Guide
- Database Administrator's Guide (Advanced)
- Glossary of Terms
- Installation Guide
- Product Overview
- Quick Start
- Release Notes
- SQL Programmer's Guide
- SQL Reference Manual
- Troubleshooting Guide

# Suggested Reading Paths

---

This section provides a suggested reading path for various types of users. Read the manuals listed under All Users first.

## All Users

- Product Overview (basic concepts critical to understanding Vertica)
- Quick Start (step-by-step guide to getting Vertica up and running)
- Glossary of Terms (glossary of terms)

## System Administrators

- Installation Guide (platform configuration and software installation)
- Release Notes (release-specific information)
- Troubleshooting Guide (general troubleshooting information)

## Database Administrators

- Installation Guide (platform configuration and software installation)
- Database Administrator's Guide (database configuration, loading, security, and maintenance)
- Troubleshooting Guide (general troubleshooting information)

## Application Developers

- SQL Programmer's Guide (connecting to a database, queries, transactions, etc.)
- SQL Reference Manual (Vertica-specific language information)
- Troubleshooting Guide (general troubleshooting information)





## Where to Find Additional Information

Visit the **Vertica Systems, Inc. Web site** (<http://www.vertica.com>) to keep up to date with:

- Downloads
- Frequently Asked Questions (FAQs)
- Discussion forums
- News, tips, and techniques

## Typographical Conventions

It is important to understand the terms and typographical conventions used in this document.

General Convention	Description
<b>colored bold text</b>	introduces new terms defined either in the text, the glossary, or both.
<i>normal italic text</i>	indicates emphasis and the titles of other documents.
UPPERCASE TEXT	indicates the name of an SQL command or keyword.
monospace text	indicates literal interactive or programmatic input/output.
<i>italic monospace text</i>	indicates user-supplied information in interactive or programmatic input/output.
<b>bold monospace text</b>	indicates literal interactive user input
↵	indicates the Return/Enter key; implicit on all user input that includes text
SQL Syntax Convention	Description
indentation	is an attempt to maximize readability; SQL is a free-form language.
backslash \	continuation character used to indicate text that is too long to fit on a single line.
braces { }	indicate required items.
brackets [ ]	indicate optional items.
ellipses ...	indicate an optional sequence of similar items.
vertical ellipses ≡	indicate an optional sequence of similar items or that part of the text has been omitted for readability.
vertical line	within braces or brackets, indicates a choice .

# Preface

---

This document describes how to solve problems that you might encounter in setting up and maintaining a Vertica database.

## Prerequisites

This document assumes that you have become familiar with the concepts discussed in the Product Overview.

Vertica recommends that you follow the tutorial in the Quick Start to get experience with the process of setting up an example database before you configure a production database.

## Audience

This document is intended for anyone with responsibility for configuring, loading, securing, and maintaining a Vertica database.



# Overview

---

Most Vertica problems are the result of Linux or network configuration errors. Begin by make sure that all steps listed the following sections of the Installation Guide have been carried out:

1. Linux Configuration
2. Network Configuration

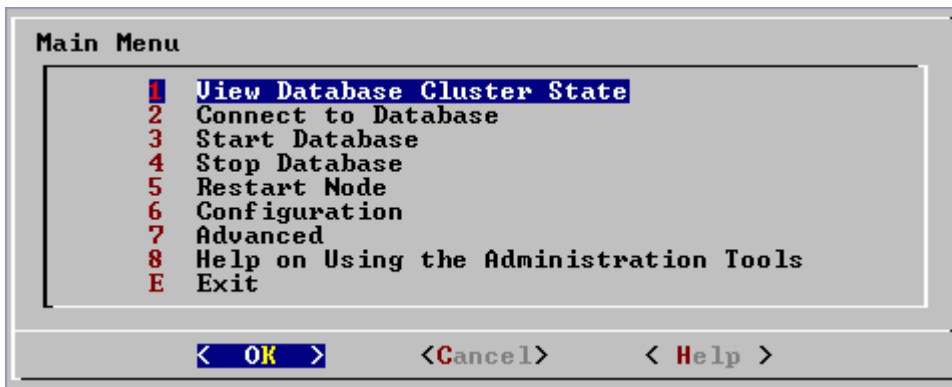
It is essential to Check Hostname Resolution as described in the Installation Guide. Invalid hostname resolution is a very common source of problems. Your Vertica database will not work unless the hostname resolution is correct.

# Monitoring the Database

## View Database Cluster State

The Administration Tools provide a View Database Cluster State tool. This tool shows the current state of the nodes in the database.

1. On the Main Menu, select **View Database Cluster State**.



2. The normal state of a running database is ALL UP. The normal state of a stopped database is ALL DOWN.

DB _____	Node	State
Clickstream_Schema_____	ALL_	DOWN_
CreditHistory_Schema_	ALL_	DOWN_
Retail_Schema_____	ALL_	DOWN_
Stock_Schema_____	ALL_	UP_
Telecom_Schema_____	ALL_	DOWN_

< OK >

- If some nodes are UP and some DOWN, restart the database as described in Starting and Stopping the Database (unless you have a known node failure and wish to continue in that state.)

DB	Node	State
Clickstream_Schema	ALL	DOWN
CreditHistory_Schema	ALL	DOWN
Retail_Schema	ALL	DOWN
Stock_Schema	stock_schema_node1_host01	UP
Stock_Schema	stock_schema_node2_host02	UP
Stock_Schema	stock_schema_node3_host03	DOWN
Stock_Schema	stock_schema_node4_host04	UP
Telecom_Schema	ALL	DOWN

< OK >

- Nodes that are INITIALIZING or RECOVERING indicate that **failure recovery** (page 31) is in progress.
- Nodes in other states (such as NEEDS\_CATCHUP) are transitional and can be ignored unless they persist. In that case, call **Technical Support** (on page 11).

For more information, see:

- the Advanced section of the Administration Tools
- Startup Problems** (page 33)
- Shutdown Problems** (page 32)

## Monitoring the Log Files

### When a Database is Running

When a Vertica database is running, each node in the cluster writes messages into a file named **vertica.log**. For example, the Tuple Mover and the transaction manager write INFO messages into vertica.log at specific intervals even when there is no WOS activity.

To monitor a running database in real time:

- Log into the database administrator account on any or all of the nodes in the cluster.
- Enter:

```
$ tail -f catalog-path/database-name/node-name_catalog/vertica.log
```

<i>catalog-path</i>	the catalog pathname specified when you created the database (see Create Database in the Database Administrator's Guide).
<i>database-name</i>	the database name (case sensitive)

<code>node-name</code>	the node name as specified in the database definition (see View Database in the Database Administrator's Guide).
------------------------	--

### When No Database is Running

When no database is running, each node in the cluster writes messages into a file named **dbLog**. For example, if a database fails to start before it can write messages into `vertica.log`, check the file:

```
catalog-path/database-name/dbLog
```

where *catalog-path* and *database-name* are as described above.

### See Also

- **Rotating Log Files** (page 24)

## Rotating the Log Files

### Using the logrotate Tool

You can use the Linux tool **logrotate** [http://www.linuxcommand.org/man\\_pages/logrotate8.html](http://www.linuxcommand.org/man_pages/logrotate8.html) (see `man logrotate`) to manage Vertica log files. The steps to configure log rotation for your installation are as follows:

1. On each node, edit the file `/opt/vertica/config/vertica.logrotate` to provide the paths to your database log files and to alter configuration parameters as required. By default, `vertica.logrotate` operates on both `vertica.log` and `dbLog`, as described in **Monitoring the Log Files** (page 23). It rotates each log file once a week, compresses them, and keeps them on disk for one year.
2. Copy the `vertica.logrotate` file to the system `logrotate` directory and change its name to **vertica**.

```
$ cp /opt/vertica/config/vertica.logrotate /etc/logrotate.d/vertica
```
3. Optionally create a cron job to automatically rotate the log files.

### Manual Log Rotation

To perform manual log rotation, use the following procedure. You can use this procedure to implement a custom log rotation process. No log messages are lost during the procedure.

1. Rename or archive the `vertica.log` file that is produced. For example:

```
$ mv vertica.log vertica.log.1
```
2. Send the Vertica process the USR1 (gracefully restart) signal. For example:

```
$ killall -USR1 vertica
```

or



```
$ ps -ef | grep -i vertica
$ kill -USR1 process-id
```

## Using the SQL Monitoring API

Vertica provides an API for monitoring various features and functions within a database in the form of virtual tables that can be queried using a limited form of the SELECT statement. You can use external tools to query the virtual tables and act upon the information as desired. For example, you can use a third-party monitoring tool to periodically query the K-Safety level of the database. If it falls below the desired level indicating a host failure, the monitoring tool can use any means necessary to notify the database administrator and/or appropriate IT personnel.

The virtual tables that make up the monitoring API are described in the SQL Reference Manual. They are:

Virtual Table	Description
VT_COLUMN_STORAGE	monitors the amount of disk storage used by each column of each projection on each node.
VT_DISK_STORAGE	monitors the amount of disk storage used by the database on each node.
VT_LOAD_STREAMS	monitors load metrics for each load stream on each node.
VT_PROJECTION_STORAGE	monitors the amount of disk storage used by each projection on each node.
VT_QUERY_METRICS	monitors the sessions and queries executing on each node.
VT_RESOURCE_USAGE	monitors system resource management on each node.
VT_SYSTEM	monitors the overall state of the database.
VT_TABLE_STORAGE	monitors the amount of disk storage used by each table on each node.
VT_TUPLE_MOVER	monitors the status of the Tuple Mover on each node.

### Query Syntax

Virtual table queries use a different processing mechanism than database queries. Thus, the virtual tables support only a very limited set of query capabilities:

- Only the FROM Clause is allowed in the SELECT statement.
- Historical queries are not allowed, including AT EPOCH LATEST.

When a cluster is in the Recovering state, the server refuses connection requests and thus cannot be monitored using the SQL monitoring API.

## Examples

```
SELECT CURRENT_EPOCH, K_SAFETY FROM VT_SYSTEM;
SELECT * FROM VT_RESOURCE_USAGE;
SELECT NODE, TOTAL_USER_SESSIONS, TOTAL_QUERIES_EXECUTED FROM
VT_QUERY_METRICS;
```

## Monitoring Processes

You can use `ps` to monitor the database and Spread processes running on each node in the cluster. For example:

```
$ ps aux | grep /opt/Vertica/bin/Vertica
$ ps aux | grep /opt/Vertica/sbin/spread
```

You should see exactly one Vertica process and exactly one Spread process on each node. To monitor Administration Tools and connector processes:

```
$ ps aux | grep Vertica
```

There may be many connection processes but at most one Administration Tools process.

## Monitoring System Resource Usage

You may find it helpful to monitor system resource usage on any or all nodes in the cluster.

1. Log into the database administrator account on a node.
2. Run the `top` utility. A high CPU percentage in `top` indicates that Vertica is CPU-bound. For example:

```
top - 16:08:52 up 7 days, 4:52, 9 users, load average: 0.91, 0.97, 0.81
Tasks: 123 total, 1 running, 122 sleeping, 0 stopped, 0 zombie
Cpu(s): 26.9% us, 1.3% sy, 0.0% ni, 71.8% id, 0.0% wa, 0.0% hi, 0.0% si
Mem: 4053136k total, 3882020k used, 171116k free, 407688k buffers
Swap: 4192956k total, 176k used, 4192780k free, 1526436k cached
```

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
8785	dbadmin	1	0	1374m	678m	55m	S	99.9	17.1	80:41.14	vertica
3557	root	16	0	32152	11m	2508	S	1.0	0.3	53:30.44	X
1	root	16	0	4748	552	456	S	0.0	0.0	0:00.75	init
2	root	RT	0	0	0	0	S	0.0	0.0	0:00.14	migration/0
3	root	34	19	0	0	0	S	0.0	0.0	0:00.04	ksoftirqd/0
4	root	RT	0	0	0	0	S	0.0	0.0	0:00.10	migration/1
5	root	34	19	0	0	0	S	0.0	0.0	0:00.03	ksoftirqd/1
6	root	RT	0	0	0	0	S	0.0	0.0	0:00.09	migration/2
7	root	34	19	0	0	0	S	0.0	0.0	0:00.03	ksoftirqd/2

Some possible reasons for high CPU usage are:

- The Tuple Mover runs automatically and thus consumes CPU time even if there are no connections to the database. If you believe this to be a problem, contact **Technical Support** (on page 11).

- The `pdflush` process (a set of worker threads for writing back dirty filesystem data) is consuming a great deal of cpu time, possibly driving the load way up. Adding RAM appears to make the problem worse. Log in to root and change the Linux parameter `swappiness` to 0.

```
# echo 0 > /proc/sys/vm/swappiness
```

- Some information sources:

**TechRepublic** <http://techrepublic.com.com/5206-6230-0.html?forumID=36&threadID=175191&start=0>

**Red Hat** [https://bugzilla.redhat.com/bugzilla/show\\_bug.cgi?id=150653](https://bugzilla.redhat.com/bugzilla/show_bug.cgi?id=150653)

**Indiana University Unix Systems Support Group**

<http://www.ussg.iu.edu/hypermail/linux/kernel/0404.3/0744.html>

3. Run the `iostat` utility. A high idle time in `top` at the same time as a high rate of blocks read in `iostat` indicates that Vertica is disk-bound. For example:

```
$ /usr/bin/iostat
```

```
Linux 2.6.9-22.ELsmp (qa0) 07/13/2007
```

```
avg-cpu:  %user   %nice   %sys %iowait   %idle
           0.83    0.00    0.13    0.02   99.03
```

```
Device:            tps    Blk_read/s    Blk_wrtn/s    Blk_read
   Blk_wrtn
hda                 0.37         3.40         10.37        2117723
   6464640
sda                 0.46         1.94         18.96        1208130
   11816472
sdb                 0.26         1.79         15.69        1114792
   9781840
sdc                 0.24         1.80         16.06        1119304
   10010328
sdd                 0.22         1.79         15.52        1117472
   9676200
md0                 8.37         7.31         66.23        4554834
   41284840
```



# Using Diagnostic Tools

---

Vertica provide the following diagnostic tools:

- ***Diagnostics Utility*** (page 30)
- ***DUMP\_CATALOG*** (page 30)
- ***DUMP\_LOCKTABLE*** (page 31)

## Determining the Version Number

To determine which version of Vertica is installed on a host, log in to that host and type:

```
# rpm -qa | grep vertica
<package-name>
```

The name of the rpm package that is installed contains the version and build numbers.

## Diagnostics Utility

The diagnostics utility gathers up database logs along with other useful information into a single .zip file.

### Syntax

```
/opt/vertica/bin/diagnostics [ command... ]
```

### Semantics

<i>command</i>	-h --help	shows a help message and exits
	-l --listdbs	lists the running databases
	-s --local_diags	gather diagnostics for local machine only
	-d <i>dbname</i> --database <i>dbname</i>	retrieves information about the database <i>dbname</i>

### Notes

- In most cases, running `/opt/vertica/bin/diagnostics` without any arguments should be sufficient; it reports on all databases and all nodes.
- The diagnostics utility writes the results into a file named `/tmp/VerticaDiagnostics.date.zip` where *date* is replaced with a numeric date representation so that multiple calls to diagnostics result in different files being created.

## DUMP\_CATALOG

If Vertica seems to be hung, use DUMP\_CATALOG to check for deadlocked clients and the resources they are waiting for:

1. Open an additional vsql connection.
2. Execute the query:  

```
SELECT DUMP_CATALOG();
```

The output is written to vertica.log (see **Monitoring the Log Files** (page 23)),
3. Send the output to **Technical Support** (on page 11).

## DUMP\_LOCKTABLE

1. Open an additional vsql connection.

2. Execute the query:

```
SELECT DUMP_LOCKTABLE( );
```

The output is written to vertica.log (see *Monitoring the Log Files* (page 23)),

3. Send the output to **Technical Support** (on page 11).

## Failure Recovery

Recovery is the process of restoring the database to a fully-functional state after one or more nodes in the system has experienced a software or hardware related failure. Vertica has a unique approach to recovering a node that is based on querying replicas of the data stored on other nodes. For example, a hardware failure may cause a node to lose database objects or to miss changes made to the database (INSERTs, UPDATEs, etc.) while offline. When the node comes back on line, it recovers lost objects and catches up with changes by querying the other nodes.

K represents the maximum number of nodes in a database that can fail and recover with no loss of data. In Vertica V2.0, the value of K can be zero (0) or one (1). The value of K can be one (1) only when the Physical Schema design meets certain requirements. The designs generated by the Database Designer are K-Safe.

For a database to be safe, no more than K nodes can be down, as shown below.

Nodes	Nodes DOWN	State of Database
4	0	Safe.
4	1	Precarious. Data loss can occur.
4	2	Inoperative. Automatic shutdown.

### Automatic Recovery

Automatic recovery handles startup after a single node failure without intervention. For example, when a failed node comes back on line and rejoins the database, it recovers its lost data by querying the other nodes, as long as there are enough active nodes to ensure K-Safety. Transactions can continue to commit during the recovery process, except for a short period at the end of the recovery.

In the case of multiple node failures, automatic recovery shuts down the database. When you restart the database, automatic recovery attempts to start all the nodes and to recover objects on any node that lost objects as a result of a failure. If successful, View Database Cluster State shows all nodes UP. Otherwise, automatic recovery writes out information that can be used for manual recovery and then shuts down again.

### Manual Recovery

When automatic recovery fails to restart the database, use the Start Database command and follow the built-in manual recovery steps. The database is not available for connections until manual recovery is complete.

### Administrator Specified Recovery

Situations not handled by the built-in manual recovery steps are documented in both the Database Administrator's Guide and the Troubleshooting Guide for convenience. These are classified as:

- **Shutdown Problems** (page 32)
- **Startup Problems** (page 33)
- Node Failures

Tools for correcting these situations are described in the Advanced section of the Administration Tools.

## Shutdown Problems

This section describes some of the known problems that can occur when stopping a database. Vertica automatically shuts down a database when one of the following events occurs:

- the administrator uses the Stop Database command.
- the cluster becomes unsafe.

### Cannot shut down while users are connected

#### Error Message

```
NOTICE:  Cannot shut down while users are connected
```

#### Explanation

There are still open connections to the database.

#### Workaround

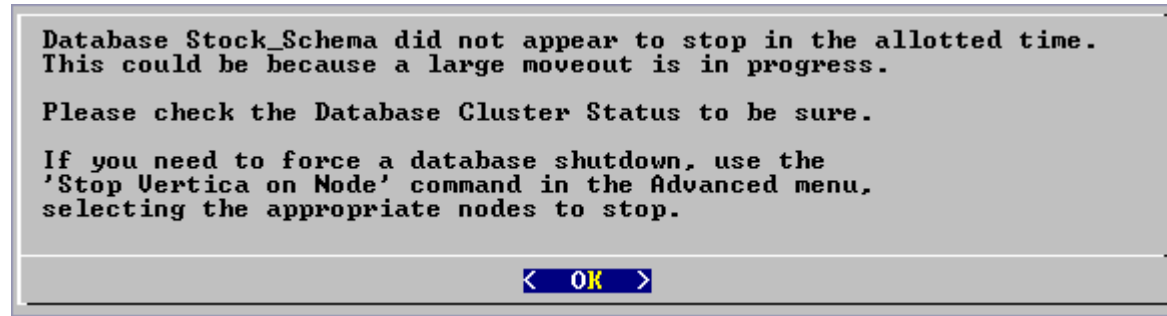
Make sure that there are **no open connections** to the database. On each host:



```
$ netstat | grep 5433
tcp        0      0 host01:5433          host01:33558        ESTABLISHED
tcp        0      0 host01:33558        host01:5433        ESTABLISHED
```

## Database ... did not appear to stop...

### Error Message



### Explanation

If you have a large database, the Vertica process may need more time to complete a moveout. Shutting down the database cancels in-progress sessions and stops the ATM so it is unlikely that any other operations have not completed.

### Workaround

Vertica Systems, Inc. recommends that you wait as long as possible before taking action. You can cause data loss by, for example, interrupting a database that is still performing a moveout.

1. **Check the log files** (page 23) to see if any messages are being logged. A moveout causes a pair of <INFO> messages to appear in the log files. If you a message indicating that a moveout has begun but no message indicating that it finished (committed), the moveout is still running.
2. Check the **processes** (page 26) and **system resource usage** (page 26) for database activity.
3. If you are certain that no moveouts are still in progress, go to **Advanced** > Stop Vertica on Node.
4. If that does not work, go to **Advanced** > Kill Vertica Process on Node. This command forces the cluster to go through recovery at startup.

## Startup Problems

This section describes some of the known problems that can occur when starting a database. Startup fails on:

- individual nodes when automatic recovery fails for a node.

- all nodes when:
  - automatic recovery is not possible
  - manual recovery was not specified properly
  - manual recovery failed
  - the cluster becomes unsafe

If a database fails to start before it can write messages into `vertica.log`, check the file `catalog-path/database-name/dbLog`.

## Startup successful, but some nodes are recovering

### Error Message

```
Startup successful, but some nodes are recovering. You can use the View
Database Cluster State option to check progress.
Press RETURN to continue
```

### Explanation

This message typically indicates an abnormal shutdown of one or more nodes.

### Workaround

Postpone query and load processing until all nodes are up.

## Error starting database, no nodes are up

### Error Message

```
Error starting database, no nodes are up
Press RETURN to continue
```

### Explanation

An unknown problem is preventing the database from starting.

### Workaround

Reboot the hosts and try to start the database. If unsuccessful, contact **Technical Support** (on page 11).

## Database startup successful, but it may be incomplete

### Error Message

Database startup successful, but it may be incomplete. Some nodes remain in a transitional state. This state may be caused by cluster catalog inconsistency. See Database Cluster State in Main Menu for details. If this state persists, try using the Advanced menu to Stop Vertica on all nodes, then Restart Database with Catalog Version.  
Press RETURN to continue

### Explanation

Some nodes are in a transitional state: not up but not recovering.

### Workaround

If this error persists, try using the Stop Vertica on Node command in the Advanced menu to stop Vertica on all nodes. Then use the Roll Back Database With Catalog Version command.

### Error Message

Database startup successful, but it may be incomplete. Some nodes remain in a transitional state.  
See Database Cluster State in Main Menu for details.  
Press RETURN to continue

### Explanation

This message indicates that the Database Administrator chose not to wait when prompted and that the database cannot start.

### Workaround

If this error persists, contact **Technical Support** (on page 11).

## Database did not start cleanly on initiator node!

### Error Message

```
ERROR: Database did not start cleanly on initiator node!  
Stopping all nodes  
Issuing shutdown command to database
```

### Explanation

Configuration problems can cause this error.

## Workaround

1. Check hostname resolution as described in Check Hostname Resolution section of the Installation Guide.
2. Examine `/etc/hosts` on each node and specify a fully qualified domain name and an unqualified hostname. For example:  
`192.168.1.99 node01.fqdomain.com node01`

## TIMEOUT ERROR: Could not login with SSH

### Error Message

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

### Explanation

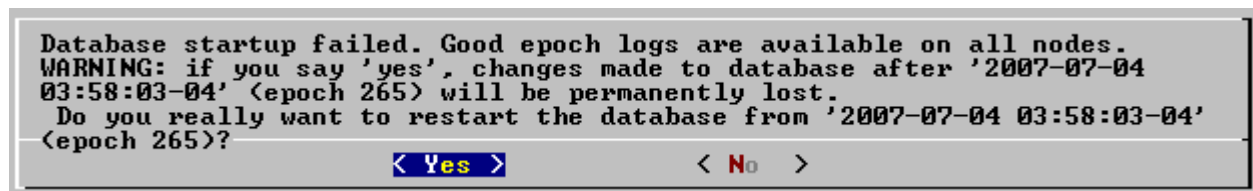
Installing Vertica on a host that is missing the mount point `/dev/pts` may result in the error when creating a database.

### Workaround

Make sure that `/dev/pts` is mounted..

## Good epoch logs are available on all nodes

### Error Message



```
Database startup failed. Good epoch logs are available on all nodes.  
WARNING: if you say 'yes', changes made to database after '2007-07-04  
03:58:03-04' (epoch 265) will be permanently lost.  
Do you really want to restart the database from '2007-07-04 03:58:03-04'  
(epoch 265)?  
< Yes > < No >
```

### Explanation

A startup attempt failed due to database inconsistency across the cluster. Vertica has determined that it can probably restart and continue running at an earlier epoch.

### Workaround

Restarting from the suggested epoch erases any changes made to the database subsequent to that epoch, across the cluster. It is likely that these changes were incomplete and erasing them will allow the cluster to proceed normally using the data saved prior to the epoch.

## No good epoch log available on node

### Error Message



### Explanation

There are a number of possible reasons for this error message, including an abnormal startup or shutdown. Every node in the cluster must be started with the same recovery epoch. Non-matching recovery epochs occur when a cluster has experienced an unsafe shutdown.

### Workaround

1. Make sure that all nodes are powered on.
2. Start the database again.
3. Make sure that all nodes have Spread running (see Check Spread). If necessary, restart Spread where it is not running and start the database.
4. On each node that did not start up, examine **dbLog** (page 22) for the cause of the failure.
5. If the cause cannot be determined, it's likely that a node has no catalog version or epoch log from which to recover. Run diagnostic tests (see **Using Diagnostic Tools** (page 29)) and contact **Technical Support** (on page 11).

## Nodes stuck in INITIALIZING state

### Error Message

In rare cases, some or all nodes can get stuck in the INITIALIZING state when trying to start the database.

### Explanation

This is known to happen when a database crashes while CREATE TABLE or other DDL commands are executing.

### Workaround

1. **Advanced** > Stop Vertica on Node to stop all nodes.
2. **Main Menu** > Start Database.

## Spread is not running

### Error Message

```
spreadCheck fails
spread may not be running on all nodes
    Error while starting/enabling multicasting to all hosts
    The following hosts are not running spread:
```

### Workaround

1. Use Check Spread to verify that Spread is not running.
2. Examine `/tmp/adminTools-username.log` for problems.
3. Examine `/tmp/spread*.log` and `/var/log/spreadd.log` for problems.
4. Make sure the `/etc/hosts` file is correct (see Define the Loopback Address in the Installation Guide).
5. If necessary, restart Spread.

### Restart the Spread Daemon

1. Log in as root.  

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

You can use `sudo` (if enabled) if you do not have the root password.
2. Restart the Spread daemon:  

```
# /etc/init.d/spreadd restart
```
3. Make sure the daemon is running:  

```
# ps ax | grep spread
```

# Installation Problems

---

## Failed Dependencies

### Error Message

error: Failed dependencies:

### Explanation

One or more the software products required by the Vertica rpm package is missing.

### Workaround

Install the missing product. For example:

### Readline

The Readline run-time library provides a set of functions that allow users to:

- edit command lines
- maintain a list of previously-entered command lines
- perform csh-like history expansion on previous commands.

To install libreadline:

1. Download libreadline.so.4:
  - **32-bit version** [http://mirrors.kernel.org/fedora/core/3/x86\\_64/os/Fedora/RPMS/readline-4.3-13.i386.rpm](http://mirrors.kernel.org/fedora/core/3/x86_64/os/Fedora/RPMS/readline-4.3-13.i386.rpm)
  - **64-bit version** [http://mirrors.kernel.org/fedora/core/3/x86\\_64/os/Fedora/RPMS/readline-4.3-13.x86\\_64.rpm](http://mirrors.kernel.org/fedora/core/3/x86_64/os/Fedora/RPMS/readline-4.3-13.x86_64.rpm)
2. Use rpm to install the package. For example:  
\$ `rpm -Uvh readline-version-platform.rpm`  
where *version-platform* specifies the name of the package you downloaded.

### Python

A Python interpreter must be installed on any host on which you want to use the Administration Tools.

Red Hat and Fedora Linux automatically install Python.

If you do not already have Python installed on all hosts, you can download python from the <http://www.vertica.com> (<http://www.vertica.com>).

Python is also installed by host-config.pl. It can't be right to install it twice. When is the right time? Before or after the rpm installation?

## SSH already configured

### Error Message

```
SSH already configured
If root does not have password-less you will be prompted three times for each host
```

### Explanation

This error occurs when using `install_vertica`. It indicates that ssh is already configured for root so that each connection requires a password.

### Workaround

It is possible that `dba-ssh-config` succeeded. Log in to the administrative account and test ssh by connecting to each other host. If unsuccessful, log in to the console on each host and perform the Enable Secure Shell (SSH) Logins procedure in the Installation Guide.



# Configuration Problems

---

## Administration Tools Errors

### Explanation

The Vertica administration tools expect the user's prompt to have standard bash type characters. If the prompt does not match the regular expression string is "[#\$]~[#\$]bash.\*?[#\$][#\$] "). the administration tools fails and blame SSH.

### Workaround

Change the shell prompt to use standard bash type characters.

## Cannot find IP address for host...

### Error Message

```
Cannot find IP address for host...
```

### Explanation

If you encounter this error when creating a new database, `/etc/hosts` may not be configured properly.

### Workaround

1. See Check Hostname Resolution in the Installation Guide.
2. Do not use the loopback address or associated hostname.
3. Check `/opt/vertica/config/dbname.conf` for unknown. This means Vertica cannot resolve the hostname.

## Clock skew detected

### Error Message

```
WARNING Clock skew detected; this node is at least nnn seconds fast
```

### Explanation

Vertica has detected that one or more hosts require system clock adjustment.

## Workaround

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.

For instructions, see:

- [http://kbase.redhat.com/faq/FAQ\\_43\\_755.shtm](http://kbase.redhat.com/faq/FAQ_43_755.shtm)  
http://kbase.redhat.com/faq/FAQ\_43\_755.shtm
- [http://kbase.redhat.com/faq/FAQ\\_43\\_2790.shtm](http://kbase.redhat.com/faq/FAQ_43_2790.shtm)  
http://kbase.redhat.com/faq/FAQ\_43\_2790.shtm

## Could not login with SSH

### Error Message

```
EOF ERROR: Could not login with SSH. Here is what SSH said:  
ssh: \033[D\033[Ctest01_node02: Name or service not known
```

### Explanation

If you get this message when entering a hostname for a node in the Administration Tools (defining a cluster), the hostname may not be in the list of known hosts.

### Workaround

Log in to the host from the command line:

```
> ssh node01  
Warning: Permanently added '...' (RSA) to the list of known hosts.  
Last login: Thu Feb 22 21:29:31 2007 from ...  
> exit
```

If the workaround does not solve the problem, check your SSH configuration as described in Enable Secure Shell (SSH) Logins in the Installation Guide.

## Could not open configuration file

### Error Message

```
vsq1: FATAL: could not open configuration file "<filename>": No such file or  
directory
```

### Workaround

Make sure that PostgreSQL is not installed on any host.

## Create Database Fails

### Error Message

```
create DB failed due to node connectivity issues
```

### Explanation

The first time you connect, ssh asks if you want to accept the unknown fingerprint from the new node.

### Workaround

When setting up passwordless ssh, connect to each node once to establish the fingerprint. See [dba-ssh-config](#) and [Enable Secure Shell \(SSH\) Logins](#) in the Installation Guide for more information.



## Database Designer Errors

### No such file or directory ... super\_projections.tmp.sql

#### Error Message

```
No such file or directory:  
'/tmp/vertica_test_new/super_projections.tmp.sql'
```

#### Workaround

Create the directory manually. For example:

```
> mkdir /tmp/vertica_test_new
```

### Error detected during node verification

#### Error Message

```
Error detected during node verification, host: hostname  
Cannot create database on hostname
```

#### Workaround

Check the Check Hostname Resolution section of the Installation Guide.

### Not enough open file handles allowed...

#### Explanation and Workaround

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` and adding the lines:  

```
* soft nofile 65536  
* hard nofile 65536
```

The leading asterisk means that the setting applies to all users.
3. Reload the config file:  

```
# sysctl -p
```
4. Log out and log in again. The new shell will have the new ulimit setting.

### Note

If the `nofile` parameter is less than 32768, Vertica fails and logs the following message:

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

## Permission denied

### Error Message

```
I/O error(13): Permission denied
```

### Explanation

The configuration directory is writable only by root.

### Workaround

1. Log in as root:

```
$ su - root
password: <root-password>
#
```
2. Change the permissions on the configuration directory:

```
# chmod a+rw /opt/Vertica/config
```

## Spread Panic

### Error Message

```
Error while starting/enabling multicasting to all hosts
Spread panic during re-init on the following hosts: ['vertica01']
```

### Explanation

Vertica automatically sets up a Spread configuration for the cluster when you use the Create Database command in the Administration Tools, and starts Spread. Various other configuration errors can cause the Spread startup to fail.

### Workaround

Do not attempt to change the Spread configuration. Contact **Technical Support** (on page 11).

# Connection Problems

---

## could not connect to database

### Error Message

```
could not connect to database
connection refused
server closed the connection unexpectedly
```

### Explanation

You probably have hostname resolution problems.

### Workaround

Check `/opt/vertica/config/dbname.conf` for unknown. Make sure you have followed the instructions in the Check Hostname Resolution section of the Installation Guide.

## could not connect to server: connection refused

### Error Message

```
could not connect to server: Connection refused
Is the server running locally and accepting
connections on socket "/tmp/.s.VTPORT.5433"?
```

### Explanation

A stale socket exists in `/tmp`. This usually happens when the db is not shut down normally.

### Workaround

Remove `/tmp/.s.VTPORT.5433*` and restart the database.

## Firewall Configuration

Vertica does not recommend using firewalls between nodes. A database cluster should depend on an external firewall provided by a gateway, not Linux. For optimal performance, the host in the cluster machines must work together as though they were a single powerful machine. Firewalls prevent this from happening.

If you must use firewalls, open the following ports on each.

**Vertica**

5433 TCP (All connections)

**Spread**

4803 TCP (Client connections)

4803 UDP (Daemon <-> Daemon)

4804 UDP (Daemon <-> Daemon)

4805 UDP (Monitor to Daemon) (optional and only if "DangerousMonitor = yes" in config file)



# SQL Execution Problems

---

## Column Definition Changed to NOT NULL

### Error Message

WARNING: column "... " definition changed to NOT NULL

### Explanation

In a CREATE TABLE statement, you must specify NOT NULL on columns that will be given PRIMARY and FOREIGN KEY constraints. Otherwise, Vertica sets the column to NOT NULL and issues a warning.

### Workaround

Always specify NOT NULL on columns that will be given PRIMARY and FOREIGN KEY constraints.

## CREATE PROJECTION Hangs

### Explanation

In theory, this should never occur.

### Workaround

If the CREATE PROJECTION statement is still hanging after five minutes:

1. Open an additional vsql connection.
2. Execute the query:  

```
SELECT DUMP_LOCKTABLE();
```

The output is written to vertica.log (see *Monitoring the Log Files* (page 23)),
3. Send the output to *Technical Support* (on page 11).

## Ctrl-C Distorts Query Result Set Output

### Explanation

If you interrupt an interactive query by typing Ctrl-c after the result set starts to display, your terminal settings may adversely affected. This is caused by vsql sending its output to an external pager program. For more information about pagers, see:

- -P assignment --pset assignment
- \pset parameter [ value ]

### Workaround

- \pset pager — enables and disables (toggles) the pager program.
- \!reset — resets (clears) the query buffer.

## DEPRECATED syntax

### Explanation

When using range segmentation in Vertica V2.0 or later, the segmentation column must be a table-column, not a projection-column. Otherwise, you get a message similar to:

```
WARNING: DEPRECATED syntax. Segment expression "c1" is a projection column name, segmenting on attribute "pk1" of table "dim_tab" instead
```

### Workaround

No workaround is needed. The message can be ignored unless you intend to use the CREATE PROJECTION statement again. In that case, change c1 to pk1.

## INSERT Fails

### Workaround

You must create projections that provide storage for all columns of all tables before you can insert data. See The Physical Schema (Projections) for more information.

## Insufficient Projections to answer query

### Error Message

```
Insufficient Projections to answer query
```

## Explanation

Possible causes:

- You forgot to create projections before attempting to load data.
- There is no projection that has all of the columns needed by a query.
- A node is down and you have not defined a K=1 physical schema.

## Workaround

- Create projections before loading data.
- If you are using custom projections, make sure that you have met all the superprojection requirements listed in Defining Custom Projections.
- Make sure that all nodes are running.

## LIKE Predicate Referred to as ~~

### Error Message

```
ERROR: Operator ~~ (1209) ...
```

### Explanation

Error messages caused by the LIKE-predicate may refer to it as ~~ (two consecutive tilde characters).

### Workaround

None.

## Missing FROM-clause entry for table

### Explanation

An alias is specified for a table but not used in the WHERE clause. Example:

```
SELECT *
FROM   PRODUCT_DIMENSION P
WHERE  PRODUCT_DIMENSION.WEIGHT < 10;
```

### Workaround

Do not mix aliases and table names in the same query. Correct examples:

```
SELECT *
FROM   PRODUCT_DIMENSION P
WHERE  P.WEIGHT < 10;
```

```
SELECT *  
FROM PRODUCT_DIMENSION  
WHERE PRODUCT_DIMENSION.WEIGHT < 10;
```

## Not a Star or Snowflake Query

### Error Message

Not a Star or Snowflake Query

### Explanation

- The query is attempting to do an invalid join. For example, joining the fact table to a specific dimension table more than once is invalid.

### Workaround

- Use temporary tables to perform multiple joins to the same dimension table.

### Explanation

- The logical schema is missing primary key and/or foreign key constraints.

### Workaround

- See Adding Join Constraints in the Database Administrator's Guide.

## Out of Memory

### Explanation

Queries that return extremely large result sets can cause a memory error to occur.

### Workaround

Limit your queries to the specific columns you need. In other words, do not use SELECT \* on large tables.

## Query Performance

### Explanation

By default, queries do not use snapshot isolation and thus hold locks on the projections associated with the tables in the SELECT list. This blocks other database operations (concurrent queries) until the transaction ends.

### Workaround

Unless you need the result set of your query to include data from the current epoch, use:

```
AT EPOCH LATEST SELECT ...
```

If you need the result set of your query to include data from the current epoch, end the current transaction as soon as possible. This can be done by executing COMMIT or ROLLBACK or terminating the connection.

### Explanation

Because Vertica does not support cursors, a query that produces a large result set can consume a great deal of client resources.

### Workaround

- Add predicates to the query to reduce the size of the result set.
- Use the LIMIT clause (described in the SQL Reference Manual) to reduce the size of the result set.

## Site <node-name> got error: Too many groups.

### Error Message

```
ERROR: Site node0 got error: Too many groups. See log for more errors.
```

### Explanation

Vertica transforms DISTINCT into GROUP BY, which must fit in memory. The error indicates that there is not enough free memory.

### Workaround

Add memory to your cluster.

## Unexpected Query Results

### Workaround

Submit a problem report to *Technical Support* (on page 11) with the schema and queries that you are using as well as a brief description of the data set and expected results.

## Update and Delete Performance

### Explanation

Updates and deletes of large numbers of rows in a large WOS can be slow.

### Workaround

Commit the data so that it can be moved to the ROS, where delete and update processing is faster.

## WOS Overload

### Explanation and Workaround

The WOS can expand to hold up to 1GB of data at any time. If you try to load more data than the WOS can hold before the automatic tuple mover can advance the epoch and perform a moveout, the COPY statement fails and the current transaction rolls back. WOS overload is most likely to occur when loading a very large fact table. If you encounter this problem, consider one or more of the following:

- Load the data to disk (ROS) instead of memory (WOS).
- Split the fact table load file into multiple pieces and do multiple loads in sequence.

Refer to Advanced Storage Control in the Database Administrator's Guide (Advanced).

# Platform Problems

---

## Data Disk Spontaneously Remounts

### Explanation

A temporary data storage disk using an ext3 file system may spontaneously dismount and remount as read-only under heavy load. This is due to a known configuration problem in Red Hat Enterprise Linux 4.

### Workaround

1. Add the following line to `/etc/sysctl.conf` (the kernel sysctl configuration file):  
`vm.min_free_kbytes = 4096`
2. At the root command line prompt, enter:  
`# echo 4096 > /proc/sys/vm/min_free_kbytes`





# Index

---

## /

/etc/hosts • 39

## A

About the Documentation • 11  
Administration Tools Errors • 39

## C

Cannot find IP address for host... • 39  
Cannot shut down while users are connected • 30  
Catalog • 27  
    Dumping • 27  
Clock skew detected • 39  
Column Definition Changed to NOT NULL • 47  
Configuration Problems • 39  
Connection Problems • 45  
Copyright Notice • ii  
could not connect to database • 45  
could not connect to server  
    connection refused • 45  
Could not login with SSH • 40  
Could not open configuration file • 40  
Create Database Fails • 41  
CREATE PROJECTION Hangs • 47  
CRITICAL • 21  
Ctrl-C Distorts Query Result Set Output • 48

## D

Data Disk Spontaneously Remounts • 53  
Database ... did not appear to stop... • 31  
Database Designer Errors • 43  
Database did not start cleanly on initiator node! • 33  
Database startup successful, but it may be incomplete • 33  
dbLog • 20  
DEBUG • 21  
DEPRECATED syntax • 48  
Determining the Version Number • 28  
Diagnostics Utility • 9, 27, 28  
DISABLE • 21  
DUMP\_CATALOG • 27, 28  
DUMP\_LOCKTABLE • 27, 29, 47

## E

Error detected during node verification • 43  
Error starting database, no nodes are up • 32

## F

Failed Dependencies • 37  
Failure Recovery • 21, 29  
Firewall Configuration • 45

## G

Good epoch logs are available on all nodes • 34

## I

INFO • 21  
INSERT Fails • 48  
Installation Problems • 37  
Insufficient Projections to answer query • 48

## L

LIKE Predicate Referred to as ~~ • 49  
Lock table • 27  
Log files • 20

## M

Missing FROM-clause entry for table • 49  
Monitoring Processes • 24, 31  
Monitoring System Resource Usage • 24, 31  
Monitoring the Database • 20, 35  
Monitoring the Log Files • 21, 22, 28, 29, 31, 47

## N

No good epoch log available on node • 35  
No such file or directory ...  
    super\_projections.tmp.sql • 43  
Nodes stuck in INITIALIZING state • 35  
Not a Star or Snowflake Query • 50  
Not enough open file handles allowed... • 43

## O

Out of Memory • 50  
Overview • 19

## P

Permission denied • 44  
Platform Problems • 53  
Preface • 17  
Printing the PDF Files • 12

## Q

Query Performance • 51

## R

Reading the HTML Files • 12

Rotating the Log Files • 22

## S

Shutdown Problems • 21, 30

Site <node-name> got error

    Too many groups. • 51

Spread is not running • 36

Spread Panic • 44

SQL Execution Problems • 47

SSH already configured • 38

Startup Problems • 21, 30, 32

Startup successful, but some nodes are  
    recovering • 32

Suggested Reading Paths • 13

## T

Technical Support • 9, 12, 21, 25, 28, 29, 33, 35,  
    44, 47, 52

TIMEOUT ERROR

    Could not login with SSH • 34

TIMING • 21

TRACE • 21

Typographical Conventions • 16

## U

Unexpected Query Results • 52

Update and Delete Performance • 52

Using Diagnostic Tools • 27, 35

Using the SQL Monitoring API • 23

## V

View Database Cluster State • 20

## W

WARNING • 21

Where to Find Additional Information • 15

Where to Find the Vertica Documentation • 11

WOS Overload • 52