

Vertica® Analytic Database 4.1, Revision 1

Troubleshooting Guide

Copyright© 2006-2011 Vertica Systems, Inc.

Date of Publication: January 7, 2011



CONFIDENTIAL

Contents

Technical Support	1
--------------------------	----------

About the Documentation	2
--------------------------------	----------

Where to Find the Vertica Documentation	2
Reading the Online Documentation	2
Printing Full Books	4
Suggested Reading Paths	4
Where to Find Additional Information	6
Typographical Conventions	7

Preface	9
----------------	----------

Troubleshooting Overview	10
---------------------------------	-----------

Profiling Database Performance	11
---------------------------------------	-----------

Enabling and Disabling Profiling.....	11
Enabling Profiling for the Current Session.....	11
Disabling Profiling for the Current Session.....	12
Enabling Profiling for All Sessions	12
Disabling Profiling for All Sessions	12

Profiling a Single Statement	13
Real-time Profiling	13
Viewing Profiling Data	15
Viewing Real-time Profiling Data	16
Clearing Profiling Data	16
Collecting Query Information	18
Enabling and Disabling Query Repository	18
Monitoring the QUERY_REPO Table	18
Configuring Query Repository	23
Managing And Viewing Query Repository	24
Using Diagnostic Tools	26
Determining the Version Number	27
Diagnostics Utility	27
Exporting a Catalog	28
Failure Recovery	29
Shutdown Problems	31
... large moveout is in progress	31
... users are connected	32
No running statement, that session is idle.....	35
Startup Problems	35
Startup successful, but some nodes are recovering.....	35
Error starting database, no nodes are up	36
Database startup successful, but it could be incomplete	36
Database did not start cleanly on initiator node!	37
TIMEOUT ERROR: Could not login with SSH	37
Good epoch logs are available on all nodes.....	37
No good epoch log available on node.....	38
Nodes stuck in INITIALIZING state.....	38
Node does not recover because of lock timeouts	39
Spread Problems	39
Spread is not running	39
Administration Tools shows node state as UNKNOWN	40
Diagnosing spread problems	40

Installation Problems **42**

Failed Dependencies	42
SSH Already Configured	43
Running Diagnostics Utility for Failed Installation	43
Passwordless SSH Access to Other Sites Failed	43
Misleading Message Due to IP/hostname Mismatch	44
Host Key Verification Failure.....	45
Permission denied even after providing correct root password	45
Installation runs as sudoer when run using sudo root	45
Managing Large Grouped ROS Containers	46

Backup and Restore Problems **49**

Errors found in SAL or Catalog	49
Cannot bootstrap catalog on top of existing catalog	49
Restore fails if files present in catalog directory	49

Configuration Problems **51**

Administration Tools Errors	51
Cannot Allocate Memory.....	51
Cannot find IP Address for Host.....	51
Clock Skew Detected.....	52
Could Not Log in with SSH.....	52
Could Not Open Configuration File	53
Create Database Fails.....	53
Denied. Not a super user. Cannot create a schema	53
Error Detected During Node Verification.....	54
Insufficient Privilege USAGE on SCHEMA.....	54
Not Enough Open File Handles Allowed... ..	54
Using sudo NOFILE Parameters Not Saved in limits.conf.....	56
Unable to drop table. Must be Owner of Relation	57
I/O Error: Permission Denied	57
Permission Denied for Relation	57
One or More Nodes Did Not Open a Data Connection to this Node	58
Spread Dead but pid File Exists.....	58

Connection Problems **60**

Could Not Connect to Database.....	60
Could Not Connect to Server: Connection Refused	60
New Session Rejected Due to Limit	60
Firewall Configuration.....	61
Firewall inactivity timeout hit (Insufficient KEEPALIVE value)	62

Database Designer Errors **63**

No Such File or Directory ... super_projections.tmp.sql	63
---	----

Startup Problems **64**

KeyError: 'getpwuid(): uid not found: <uid>'64
 (LOSTCONTACT) Slow Database Startup when Node is Down64
 Vertica Fails to Start (verticad and pam_limits)64

Storage Location Problems **66**

Could Not Open Directory66
 Could Not Create Directory66
 Could Not Add Location (Permission)66
 Could Not Add Location (Unavailable).....67
 Usage Cannot be an Empty String67
 Usage Cannot be Changed to Data67
 Usage Cannot be Changed to Temp.....68
 Usage of Storage Location Cannot be Changed68
 Throughput Should be Greater Than (>) Zero (0)68
 Latency Should be Greater Than (>) One (1)68
 Cannot be Retired69
 DROP_LOCATION for Data Locations is not Supported.....69
 Location Cannot be Dropped.....69

Execution Problems **71**

Attempted to Create Too Many ROS Containers71
 Cannot Drop Partition with Nodes Down71
 Case-sensitive System Table Queries Return 0 Records72
 Column Definition Changed to NOT NULL72
 COPY, DELETE, UPDATE, or INSERT Hangs When Server Idle.....73
 Could Not Write To Filename73
 Ctrl-C Distorts Query Result Set Output73
 Deleting Records from Table with PK with Associated FK74
 Errors that Result in a Statement-Level Rollback.....74
 Errors that Result in a Transaction-Level Rollback74
 Execution Error on External Procedures.....74
 Function <function-name>(<data-type>, ...) does not exist.....75
 Function with specified name and parameters does not exist75
 Hash Table Out of Memory on COPY and REFRESH Operations76
 I/O Error77
 Insufficient Projections to Answer Query.....77
 LIKE Predicate Referred to as ~~77
 Loading Data While a Node is Down77
 Load Stream Could Appear Complete While Sorting in Progress78
 Lock Timeout Error Reports Incorrect Transaction Isolation Level78
 Missing FROM-clause Entry for Table78
 Operator Does Not Exist.....79
 Out of Memory79
 Partitioning Errors80
 Projection Errors81
 Projection Refresh Doesn't Start Until a Buddy Projection is Created82

Troubleshooting Guide

Optimizer Does Not Select the Best Projection	82
Query Continues to Execute on the Server When the Client is Killed.....	83
Query Required More Resources Than Initiator Resource Manager Estimated	83
Querying Monitoring Table When a Node is Recovering	84
Record Terminator for COPY When Loading Data from a Windows Client	84
Referential Integrity Load Violation.....	84
Too Many ROS Containers.....	84
Update and Delete Performance	85
Using Stream Name Parameter while Loading Data from STDIN	85
Variable Hash Join Performance	86

Platform Problems 88

Data Disk Spontaneously Remounts.....	88
Setting the Device Readahead Parameter for Improved Performance	88
Move Core File to a Specified Location	88
SUSE 10/11 Install Fails During Spread Configuration	89

SSL Problems 90

Unsafe permissions on private key file	90
Could not access server public certificate file.....	90
Could not access server private key file.....	90
Check of private key failed	91

Virtual IP Server Problems 92

Users Cannot Connect to the VIP	92
Resolving Keepalived Issues	93

Index 95

Copyright Notice 98

Technical Support

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

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

- 1 Go to <http://www.vertica.com/support> (*http://www.vertica.com/support*).
- 2 Click **My Support**.

You can also email verticahelp@vertica.com.

Before you report a problem, run the *Diagnostics Utility* (page 27) described in the Troubleshooting Guide and attach the resulting `.zip` file to your ticket.

About the Documentation

This section describes how to access and print Vertica documentation. It also includes *suggested reading paths* (page 4).

Where to Find the Vertica Documentation

You can read or download the Vertica documentation for the current release of Vertica® Analytic Database from the *Product Documentation Page* http://www.vertica.com/v-zone/product_documentation. You must be a registered user to access this page.

The documentation is available as a compressed tarball (.tar) or a zip archive (.zip) file. When you extract the file on the database server system or locally on the client, contents are placed in a /vertica41_doc/ directory.

Note: 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 of the software, refer to the documentation on your database server or client systems.

See Installing Vertica Documentation in the Installation Guide.

Reading the Online Documentation

Reading the HTML documentation files

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

The Vertica documentation is supported on the following browsers:

- Mozilla FireFox
- Internet Explorer
- Apple Safari
- Opera
- Google Chrome (server-side installations only)

The instructions that follow assume you have installed the documentation on a client or server machine.

Mozilla Firefox

- 1 Open a browser window.
- 2 Choose one of the following methods to access the documentation:
 - Select **File > Open File**, navigate to `..\HTML-WEBHELP\index.htm`, and click **Open**.
 - OR drag and drop `index.htm` into a browser window.

- OR press **CTRL+O**, navigate to `index.htm`, and click **Open**.

Internet Explorer

Use one of the following methods:

- 1 Open a browser window.
- 2 Choose one of the following methods to access the documentation:
 - Select **File > Open > Browse**, navigate to `..\HTML-WEBHELP\index.htm`, click **Open**, and click **OK**.
 - OR drag and drop `index.htm` into the browser window.
 - OR press **CTRL+O**, Browse to the file, click **Open**, and click **OK**.

Note: If a message warns you that Internet Explorer has restricted the web page from running scripts or ActiveX controls, right-click anywhere within the message and select **Allow Blocked Content**.

Apple Safari

- 1 Open a browser window.
- 2 Choose one of the following methods to access the documentation:
 - Select **File > Open File**, navigate to `..\HTML-WEBHELP\index.htm`, and click **Open**.
 - OR drag and drop `index.htm` into the browser window.
 - OR press **CTRL+O**, navigate to `index.htm`, and click **Open**.

Opera

- 1 Open a browser window.
- 2 Position your cursor in the title bar and right click > **Customize > Appearance**, click the **Toolbar** tab and select **Main Bar**.
- 3 Choose one of the following methods to access the documentation:
 - Open a browser window and click **Open**, navigate to `..\HTML-WEBHELP\index.htm`, and click **Open**.
 - OR drag and drop `index.htm` into the browser window.
 - OR press **CTRL+O**, navigate to `index.htm`, and click **Open**.

Google Chrome

Google does not support access to client-side installations of the documentation. You'll have to point to the documentation installed on a server system.

- 1 Open a browser window.
- 2 Choose one of the following methods to access the documentation:
 - In the address bar, type the location of the `index.htm` file on the server. For example:
file:///<servername>//vertica41_doc//HTML/Master/index.htm
 - OR drag and drop `index.htm` into the browser window.
 - OR press **CTRL+O**, navigate to `index.htm`, and click **Open**.

Notes

The `.tar` or `.zip` file you download contains a complete documentation set.

The documentation page of the **Downloads Web site** http://www.vertica.com/v-zone/download_vertica is updated as new versions of Vertica are released. When the version you download is no longer the most recent release, refer only to the documentation included in your RPM.

The Vertica documentation contains links to Web sites of other companies or organizations that Vertica does not own or control. If you find broken links, please let us know.

Report any script, image rendering, or text formatting problems to **Technical Support** (on page 1).

Printing Full Books

Vertica also publishes books as Adobe Acrobat™ PDF. The books are designed to be printed on standard 8½ x 11 paper using full duplex (two-sided) printing.

Note: Vertica manuals are topic driven and not meant to be read in a linear fashion. Therefore, the PDFs do not resemble the format of typical books. Each topic starts a new page, so some of the pages are very short, and there are blank pages between each topic.

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

The following list provides links to the PDFs.

- Release Notes
- Concepts Guide
- Installation Guide
- Getting Started Guide
- Administrator's Guide
- Programmer's Guide
- SQL Reference Manual
- Troubleshooting Guide

Suggested Reading Paths

This section provides a suggested reading path for various users. Vertica recommends that you read the manuals listed under All Users first.

All Users

- Release Notes — Release-specific information, including new features and behavior changes to the product and documentation
- Concepts Guide — Basic concepts critical to understanding Vertica

- Getting Started Guide — A tutorial that takes you through the process of configuring a Vertica database and running example queries
- Troubleshooting Guide — General troubleshooting information

System Administrators

- Installation Guide — Platform configuration and software installation
- Release Notes — Release-specific information, including new features and behavior changes to the product and documentation

Database Administrators

- Installation Guide — Platform configuration and software installation
- Administrator's Guide — Database configuration, loading, security, and maintenance

Application Developers

- Programmer's Guide — Connecting to a database, queries, transactions, and so on
- SQL Reference Manual — SQL and Vertica-specific language 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
- Training

Typographical Conventions

The following are the typographical and syntax conventions used in the Vertica documentation.

Typographical Convention	Description
Bold	Indicates areas of emphasis, such as a special menu command.
Button	Indicates the word is a button on the window or screen.
Code	SQL and program code displays in a monospaced (fixed-width) font.
Database objects	Names of database objects, such as tables, are shown in san-serif type.
<i>Emphasis</i>	Indicates emphasis and the titles of other documents or system files.
monospace	Indicates literal interactive or programmatic input/output.
<i>monospace italics</i>	Indicates user-supplied information in interactive or programmatic input/output.
UPPERCASE	Indicates the name of a SQL command or keyword. SQL keywords are case insensitive; <code>SELECT</code> is the same as <code>Select</code> , which is the same as <code>select</code> .
User input	Text entered by the user is shown in bold san serif type.
↵	indicates the Return/Enter key; implicit on all user input that includes text
Right-angle bracket >	Indicates a flow of events, usually from a drop-down menu.
Click	Indicates that the reader clicks options, such as menu command buttons, radio buttons, and mouse selections; for example, "Click OK to proceed."
Press	Indicates that the reader perform some action on the keyboard; for example, "Press Enter."
Syntax Convention	Description
Text without brackets/braces	Indicates content you type as shown.
< <i>Text inside angle brackets</i> >	Placeholder for which you must supply a value. The variable is usually shown in italics. See Placeholders below.
[<i>Text inside brackets</i>]	Indicates optional items; for example, <code>CREATE TABLE [schema_name.]table_name</code> The brackets indicate that the <code>schema_name</code> is optional. Do not type the square brackets.
{ <i>Text inside braces</i> }	Indicates a set of options from which you choose one; for example: <code>QUOTES { ON OFF }</code> indicates that exactly one of ON or OFF must

	be provided. You do not type the braces: QUOTES ON
Backslash \	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, <code>option[. . .]</code> means that you can enter multiple, comma-separated options. Note: Showing an 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>	Items that must be replaced with appropriate identifiers or expressions are shown in italics.
Vertical bar	Is a separator for mutually exclusive items. For example: <code>[ASC DESC]</code> Choose one or neither. You do not type the square brackets.

Preface

This guide describes how to solve problems that you might encounter with Vertica database setup and maintenance.

Prerequisites

This document assumes that you have become familiar with the concepts discussed in the Concepts Guide.

Note: Vertica recommends that you follow the tutorial in the Getting Started Guide to practice 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.

Troubleshooting Overview

Most Vertica problems are the result of Linux or network configuration errors. Begin by making sure that you have completed the steps listed in Before You Install.

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

Profiling Database Performance

To determine how the database is performing, you can profile the areas listed in the following table:

Profile Operator	Description
Session	Provides basic session parameters and lock time out data
Query	Provides general information about queries that ran, such as the query strings used and the duration of queries
EE	Provides detailed information about the execution run of each query

Enabling and Disabling Profiling

Vertica can collect profiling data for a single statement, a single session, or for all sessions on all nodes. Real-time profiling is always "on", without profiling being explicitly enabled.

This section describes how to enable and disable profiling for the current session and for all sessions.

Tip: To quickly determine if profiling is enabled, use the following command:

```
=> SELECT SHOW_PROFILING_CONFIG();

          SHOW_PROFILING_CONFIG
-----
Session Profiling: Local off, Global off
EE Profiling:      Local on, Global on
Query Profiling:   Local off, Global off
(1 row)
```

Enabling Profiling for the Current Session

To enable profiling for the current session, issue the following command:

```
=> SELECT ENABLE_PROFILING('type_of_profiling')
```

Where *type_of_profiling* is one of the following:

- `session` — Establishes profiling for sessions.
- `query` — Establishes profiling for queries.
- `ee` — Establishes profiling for query execution runs.

Examples

```
=> SELECT ENABLE_PROFILING('ee');
       ENABLE_PROFILING
-----
EE Profiling Enabled
(1 row)
```

Disabling Profiling for the Current Session

To disable profiling for the current session, issue the following command:

```
=> SELECT DISABLE_PROFILING('type_of_profiling')
```

Where *type_of_profiling* is one of the following:

- `session` — Establishes profiling for sessions.
- `query` — Establishes profiling for queries.
- `ee` — Establishes profiling for query execution runs.

Examples

```
=> SELECT DISABLE_PROFILING('ee');
       DISABLE_PROFILING
-----
EE Profiling Disabled
(1 row)
```

Enabling Profiling for All Sessions

To enable profiling for all sessions on all nodes, enter the following command:

```
=> SELECT SET_CONFIG_PARAMETER('global_profiling_type', 1)
```

Where *global_profiling_type* is one of the following

- `GlobalSessionProfiling` — Establishes profiling for sessions.
- `GlobalQueryProfiling` — Establishes profiling for queries.
- `GlobalEEProfiling` — Establishes profiling for query execution runs.

Example

The following command enables query profiling for all sessions and nodes:

```
=> SELECT SET_CONFIG_PARAMETER('GlobalQueryProfiling',1);
```

Disabling Profiling for All Sessions

To enable profiling for all sessions on all nodes, enter the following command:

```
=> SELECT SET_CONFIG_PARAMETER('global_profiling_type', 0)
```

Where *global_profiling_type* is one of the following:

- `GlobalSessionProfiling` — Establishes profiling for sessions.

- `GlobalQueryProfiling` — Establishes profiling for queries.
- `GlobalEEProfiling` — Establishes profiling for query execution runs.

Example

The following command disables query query profiling for all sessions and nodes:

```
=> SELECT SET_CONFIG_PARAMETER('GlobalQueryProfiling' ,0);
```

Profiling a Single Statement

To profile a single statement add the `PROFILE` keyword to the beginning of the statement:

```
=> PROFILE SELECT customer_name, annual_income
FROM public.customer_dimension
WHERE (customer_gender, annual_income) IN (
SELECT customer_gender, MAX(annual_income)
FROM public.customer_dimension
GROUP BY customer_gender);
```

`PROFILE < SELECT ...>` saves profiling information for future analysis.

A hint is written to `stderr` (the standard error stream and default destination for error messages and other diagnostic warnings, which are typically output to the screen) while the statement is executing:

```
NOTICE: Statement is being profiled.
HINT: select * from v_monitor.execution_engine_profiles where
transaction_id=45035996273740886 and statement_id=10;
NOTICE: Initiator memory estimate for query:
[on pool general: 1418047 KB, minimum: 192290 KB]
NOTICE: Total memory required by query: [1418047 KB]
 customer_name | annual_income
-----+-----
Meghan U. Miller |          999960
Michael T. Jackson |          999981
(2 rows)
```

Tip: Use the statement returned by the hint as a starting point for reviewing the query's profiling data.

To see what counters are available, issue the following command:

```
=> SELECT DISTINCT(counter_name) FROM EXECUTION_ENGINE_PROFILES;
```

Real-time Profiling

The real-time profiling data provides a means for monitoring long-running statements.

Real-time profiling counters are available for all currently-executing statements—including internal operations such as mergeout, recovery, and refresh—but only while the statements are executing. Unless profiling has explicitly been enabled using the keyword `PROFILE` or `enable_profiling()` (page 11), the profiling counters are not present after the statement completes.

The EXECUTION_ENGINE_PROFILES system table lets you view profiling counters, and the following list shows some of the most useful counters:

- execution time (μ s)
- rows produced
- total merge phases
- completed merge phases
- current size of temp files (bytes)

For more information about counters, see EXECUTION_ENGINE_PROFILES in the SQL Reference Manual.

Tip: In order to view real-time profiling data, you need, at a minimum, the transaction_id for the transaction of interest. If multiple statements have been executed within the transaction, then you also need the statement_id. The transaction IDs for internal operations can be found in the vertica.log files.

Example

To monitor a refresh operation that was initiated on node0001, find the “select start_refresh()” entry in the vertica.log file on node0001, which will be similar to the following log fragment:

```
2010-04-21 13:34:56.494 Refresh:0xb9ab5e0 [Refresh] <INFO> Refreshing projection public.fact_p from buddy
2010-04-21 13:34:56.494 Refresh:0xb9ab5e0 [Refresh] <INFO> Refreshing projection public.fact_p from buddy, historical epochs 0-12, oid 45035996273713164
2010-04-21 13:34:56.497 nameless:0xb972330 [Txn] <INFO> Begin Txn: a0000000000227 'Refresh through recovery'
2010-04-21 13:34:56.497 nameless:0xb972330 [Command] <INFO> Performing refresh through recovery on projection fact_p (45035996273713164) 0-12
2010-04-21 13:34:56.497 nameless:0xb972330 <LOG> @v_db_node0001: 00000: Recover alterSpec 45035996273713164 0-12
2010-04-21 13:34:56.500 nameless:0xb972330 [Command] <INFO> (a0000000000227) Executing the recovery plan
```

The transaction ID for this recovery operation is a0000000000227.

To monitor the profiling counters, you can run a command like the following:

```
=> SELECT * FROM execution_engine_profiles
WHERE TO_HEX(transaction_id)='a0000000000227'
AND counter_name = 'execution time (us)'
ORDER BY node_name, counter_value DESC;
```

In this example, find the operators with the largest execution time on each node:

```
=> SELECT node_name, operator_name, counter_value execution_time_us
FROM v_monitor.execution_engine_profiles
WHERE counter_name='execution time (us)'
ORDER BY node_name, counter_value DESC;
```

You can use the Linux watch command to monitor long-running queries with one-second updates:

```
WATCH -n 1 -d "vsq1 -c \"select node_name, operator_name, counter_value execution_time_us ... \""
```

See Also

EXECUTION_ENGINE_PROFILES in the SQL Reference Manual

Viewing Profiling Data

Vertica provides the profiling information it collects in the form of system tables that you can query using full SELECT support.

System (Virtual) Tables

The system tables for profiling are described in the SQL Reference Manual. They are as follows:

Virtual Table	Description
EXECUTION_ENGINE_PROFILES	Provides profiling information for query execution runs.
QUERY_PROFILES	Provides profiling information for queries.
SESSION_PROFILES	Provides profiling information for sessions.

Notes

The EXECUTION_ENGINE_PROFILES table contains the data for each profiling counter as a row within the table. For example the execution time (us) counter is in one row, and the rows produced counter is in a second row. Since there are many different profiling counters, many rows of profiling data exist for each operator. Some sample views are installed by default to simplify the viewing of profiling counters.

The following script creates the v_demo schema and places the views in that schema.

```
/opt/vertica/scripts/demo_eeprof_view.sql
```

There is one view for each of the profiling counters to simplify viewing of a single counter value. For example, to view the execution time for all operators, issue the following command:

```
=> SELECT * FROM v_demo.eeprof_execution_time_us;
```

To view all counter values available for all profiled queries:

```
=> SELECT * FROM v_demo.eeprof_counters;
```

To select all distinct operators available for all profiled queries:

```
=> SELECT * FROM v_demo.eeprof_operators;
```

These views can easily be combined:

```
=> SELECT * FROM v_demo.eeprof_execution_time_us
    NATURAL LEFT OUTER JOIN v_demo.eeprof_rows_produced;
```

To view the execution time and rows produced for a specific transaction and statement_id ranked by execution time on each node:

```
=> SELECT * FROM v_demo.eeprof_execution_time_us_rank
    WHERE transaction_id=45035996273709699
    AND statement_id=1 order by transaction_id, statement_id, node_name, rk;
```

To view the top five operators by execution time on each node:

```
=> SELECT * FROM v_demo.eeprof_execution_time_us_rank
    WHERE transaction_id=45035996273709699
```

```
AND statement_id=1 AND rk<=5
ORDER BY transaction_id, statement_id, node_name, rk;
```

If you see large amounts of time spent in GroupByHash operators, you might want to revisit the projection designs so that GroupByPipeline can be used. Use the Database Designer for these optimizations. (See *Creating a Physical Design* and *Creating Custom Designs* in the Administrator's Guide.)

Profiling data can also show data skews if some nodes are processing more data than others. The `rows produced` counter in the `EXECUTION_ENGINE_PROFILES` table shows how many rows have been processed by each of the operators. Comparing the `rows produced` across all nodes for a given operator would reveal if there is a data skew issue.

Note: Some of the profiling operators shown in `EXECUTION_ENGINE_PROFILES` are generic, such as Join. The EXPLAIN plan includes more details about the specific join that is executing.

Viewing Real-time Profiling Data

You can query real-time profiling data while a long-running query or other operation is executing. The `demo` view `v_demo.eeprof_execution_time_us_rank` is helpful for viewing the current top five operators by execution time prior to the query completing:

```
=> SELECT * FROM v_demo.eeprof_execution_time_us_rank
WHERE transaction_id=45035996273709699 AND statement_id=1 AND rk<=5
ORDER BY transaction_id, statement_id, node_name, rk;
```

The Linux `watch` command is useful for long-running queries or long-running internal operations by observing which part of the query plan is currently active:

```
=> watch -d -n 2 "vsq1 -c \"
=> SELECT * FROM v_demo.eeprof_execution_time_us_rank
WHERE transaction_id=45035996273709699 AND statement_id=1 AND rk<=5
ORDER BY transaction_id, statement_id, node_name, rk;\" "
```

This `watch` command executes the query every two seconds and highlights the differences between successive updates.

Tip: Using `watch` is a convenient way to monitor the currently-executing operators within the plan on each nodes in the Vertica cluster. `watch` is also a convenient way to monitor workloads that might be unbalanced between nodes; for example, if some nodes have become idle while other nodes are still active. Such imbalances could be caused by data skews or by hardware issues.

Clearing Profiling Data

Profiling data is stored in memory. Therefore, profiling can be memory intensive depending on how much data you collect.

To clear profiling data, enter the following SQL function:

```
=> SELECT CLEAR_PROFILING('type_of_profiling', 'scope')
```

Where:

- `type_of_profiling` is one of the following:
 - `session` — Establishes profiling for sessions.
 - `query` — Establishes profiling for queries.
 - `ee` — Establishes profiling for query execution runs.
- `scope` is either GLOBAL (all sessions on all nodes) or LOCAL (the current session only).

Examples

The following example clears profiling for all execution runs in the current session:

```
=> SELECT CLEAR_PROFILING('ee', 'local');
```

The following example clears query profiling for all sessions on all nodes:

```
=> SELECT CLEAR_PROFILING('query', 'global');
```

The following example clears session profiling for all sessions on all nodes:

```
=> SELECT CLEAR_PROFILING('session', 'global');
```

See Also

Collecting Query Information (page 18)

Collecting Query Information

The Vertica query repository collects information about the SQL queries that Vertica processes. You can provide this information to the Database Designer to maximize the performance of the projections it creates. You can also use the information for performance and security analysis.

The query repository analyzes queries from client sessions, as well as internal Vertica queries, such as Tuple Mover and recovery tasks to collect data. The query repository also collects data captured through query profiling. (See *Profiling Database Performance* (page 11)).

The data that the query repository collects is moved into a repository table in persistent storage. This differs from query profiling data, which is stored in memory only and is accessible through system tables.

Enabling and Disabling Query Repository

The Database Administrator can enable and disable the query repository. When enabled, the query repository collects data for all sessions. It also automatically enables query profiling on a global level if it is not already enabled. This allows query profiling to collect data that the query repository requires for all sessions.

To enable the query repository to collect data, use the following SQL command:

```
=> SELECT SET_CONFIG_PARAMETER('QueryRepositoryEnabled', '1');
```

All the data that query repository collects is stored in the QUERY_REPO table within the SYS_DBA schema.

To disable the query repository from collecting data, use the following SQL command:

```
=> SELECT SET_CONFIG_PARAMETER('QueryRepositoryEnabled', '0');
```

Note: Disabling the query repository does not disable query profiling. If you are concerned about the amount of memory used by query profiling, either disable it globally or occasionally clear the query profiling data. See *Profiling Database Performance* (page 11).

See Also

Configuration Parameters in the Administrator's Guide

Database Management Functions and QUERY_PROFILES in the SQL Reference Manual

Monitoring the QUERY_REPO Table

The data that the query repository collects is moved into a repository table in persistent storage within the SYS_DBA schema. This differs from query profiling data, which is stored in memory only and is accessible through system tables. See QUERY_PROFILES in the SQL Reference Manual.

The entire collection of data in the QUERY_REPO table consists of the following attributes:

Column Name	Data Type	Description
START_TIME	TIMESTAMP(8) NOT NULL	The time the query was started.
EPOCH_ID	INTEGER	The epoch in which the query ran.
NODE_NAME	VARCHAR(128)	The name of the node that is reporting the requested information.
USER_NAME	VARCHAR(128)	The name of the user who initiated the query.
SESSION_ID	VARCHAR(80)	The identification of the session for which query data is captured. This identifier is unique within the cluster at any point in time but can be reused when the session closes.
QUERY_IDENTIFIER	VARCHAR(65000)	A string to identify the query in system (virtual) tables. Appears in output for DML statements only. For <code>SELECT</code> column is null.
QUERY	VARCHAR(65000)	The query string used for the query.
SEARCH_PATH	VARCHAR(4000)	The contents of the search path.
DURATION	INTEGER(8)	The duration of the query in microseconds.
SCHEMA_NAME	VARCHAR(4000)	The schema name that data is loaded into. Appears in output for DML statements only. For <code>SELECT</code> column is null.
TABLE_NAME	VARCHAR(4000)	The table name that data is loaded into. Appears in output for DML statements only. For <code>SELECT</code> column is null.
PROJECTIONS	VARCHAR(4000)	The qualified projection names. This information is stored only when execution runs are profiled. See Profiling Database Performance (page 11).
EXTRA_MEMORY	INTEGER(8)	Extra memory is the amount of memory allocated to the query by the Resource Manager but not assigned to a particular operator. This is the memory from which unbounded operators pull first. If they acquire all of the extra memory, then the plan must go back to the Resource Manager for more memory.
TYPE	INTEGER(8)	Identifies the type of query: <ul style="list-style-type: none"> ▪ 0 UNKNOWN ▪ 1 SELECT ▪ 2 UPDATE ▪ 3 INSERT ▪ 4 DELETE ▪ 5 UTILITY (commands such as CREATE, COPY, and so on)

		<ul style="list-style-type: none"> ▪ 6 NOTHING (a dummy command) ▪ 7 SELECT_CATALOG (a SELECT statement from the catalog view)
ERROR_CODE	INTEGER(8)	The error code reported by Vertica.
NUM_ROWS	INTEGER(8)	The number of rows processed for the query as part of a SELECT, UPDATE, OR INSERT. (INSERT is always 1.)

Notes

- The `QUERY_REPO` table is created when you **enable the query repository** (page 18). It could take a few minutes for the table to appear unless you manually force a save; for example:


```
=> SELECT SAVE_QUERY_REPOSITORY();
      SAVE_QUERY_REPOSITORY
      -----
      Query Repository Saved
      (1 row)
```
- If the query repository has been saved to a table (`SYS_DBA.QUERY_REPO` by default) in Vertica 4.0 or prior, the table does not contain the following columns:
 - `QUERY_IDENTIFIER`
 - `EXTRA_MEMORY`
 - `SCHEMA_NAME`
 - `TABLE_NAME`
 To save these new columns, you must drop the table and recreate it to get the latest table definition. If the query repository has never been saved, or this is a new cluster, the above four columns are available in `QUERY_REPO` and no further action is required.
- In Vertica 3.5, the `QUERY_REPO` table definition changed to reflect a new maximum object length of 128 characters; thus, several of the `VARCHAR` columns increased in size. If you used the `QUERY_REPO` table before Vertica 3.5, you must drop and then recreate the table to get the latest table definition.
- Users are restricted to viewing information about their own queries, which they access using the `SYS_DBA.QUERY_REPO` view. See *Implementing Views in the Administrator's Guide*.

Example

Enable the query repository:

```
=> SELECT SET_CONFIG_PARAMETER('QueryRepositoryEnabled', '1');
```

Force a quick save on the query repository:

```
=> SELECT SAVE_QUERY_REPOSITORY();
```

Now query the `QUERY_REPO` table:

```
=> SELECT * FROM SYS_DBA.QUERY_REPO;
```

The following results are a fragment. Note that Record 37, which is for a COPY statement, shows all column output (no nulls).

```

-[ RECORD 29 ]-----+-----
START_TIME      | 2010-10-14 12:20:02.614976
EPOCH_ID       | 434
NODE_NAME      | v_vmartdb_node01
USER_NAME      | dbadmin
SESSION_ID     | raster-s1-17172:0x10
QUERY_IDENTIFIER |
QUERY          | SELECT SET_CONFIG_PARAMETER('QueryRepositoryEnabled', '1');
SEARCH_PATH    | "$user", public, v_catalog, v_monitor, v_internal
DURATION       | 2668
SCHEMA_NAME    |
TABLE_NAME     |
PROJECTIONS    |
EXTRA_MEMORY   | 0
TYPE           | 1
ERROR_CODE     | 0
NUM_ROWS       | 0
-[ RECORD 30 ]-----+-----
...
-[ RECORD 32 ]-----+-----
START_TIME      | 2010-10-14 12:25:02.012997
EPOCH_ID       | 434
NODE_NAME      | v_vmartdb_node01
USER_NAME      | dbadmin
SESSION_ID     | raster-s1-17172:0xa53
QUERY_IDENTIFIER |
QUERY          | CREATE SCHEMA sys_dba;
SEARCH_PATH    | "$user", public, v_catalog, v_monitor, v_internal
DURATION       | 6822
SCHEMA_NAME    |
TABLE_NAME     |
PROJECTIONS    |
EXTRA_MEMORY   | 0
TYPE           | 5
ERROR_CODE     | 0
NUM_ROWS       | 0
-[ RECORD 33 ]-----+-----
START_TIME      | 2010-10-14 12:25:02.030016
EPOCH_ID       | 434
NODE_NAME      | v_vmartdb_node01
USER_NAME      | dbadmin
SESSION_ID     | raster-s1-17172:0xa54
QUERY_IDENTIFIER |
QUERY          | GRANT USAGE ON SCHEMA sys_dba TO PUBLIC ;
SEARCH_PATH    | "$user", public, v_catalog, v_monitor, v_internal
DURATION       | 5906
SCHEMA_NAME    |
TABLE_NAME     |
PROJECTIONS    |
EXTRA_MEMORY   | 0
TYPE           | 5
ERROR_CODE     | 0
NUM_ROWS       | 0
-[ RECORD 34 ]-----+-----
START_TIME      | 2010-10-14 12:25:02.042181
EPOCH_ID       | 434
NODE_NAME      | v_vmartdb_node01
USER_NAME      | dbadmin
SESSION_ID     | raster-s1-17172:0xa55
QUERY_IDENTIFIER |
QUERY          | CREATE TABLE sys_dba.query_repo (START_TIME TIMESTAMP NOT NULL,EPOCH_ID INT,
          NODE_NAME VARCHAR(128),USER_NAME VARCHAR(128),SESSION_ID
          VARCHAR,QUERY_IDENTIFIER

```

Troubleshooting Guide

```

VARCHAR(65000),QUERY VARCHAR(65000),SEARCH_PATH VARCHAR(4000),DURATION INT,
SCHEMA_NAME VARCHAR(4000),TABLE_NAME VARCHAR(4000),PROJECTIONS VARCHAR(4000),
EXTRA_MEMORY INT,TYPE INT,ERROR_CODE INT,NUM_ROWS INT ) PARTITION BY cast
(START_TIME as date);
SEARCH_PATH      | "$user", public, v_catalog, v_monitor, v_internal
DURATION         | 52789
SCHEMA_NAME      |
TABLE_NAME       |
PROJECTIONS      |
EXTRA_MEMORY     | 0
TYPE             | 5
ERROR_CODE       | 0
NUM_ROWS         | 0
-[ RECORD 35 ]---
START_TIME       | 2010-10-14 12:25:02.12363
EPOCH_ID         | 434
NODE_NAME        | v_vmartdb_node01
USER_NAME        | dbadmin
SESSION_ID       | raster-s1-17172:0xa56
QUERY_IDENTIFIER |
QUERY            | CREATE VIEW sys_dba.QUERY_REPOSITORY AS SELECT * FROM sys_dba.query_repo WHERE
                    USER_NAME=CURRENT_USER;
SEARCH_PATH      | "$user", public, v_catalog, v_monitor, v_internal
DURATION         | 8625
SCHEMA_NAME      |
TABLE_NAME       |
PROJECTIONS      |
EXTRA_MEMORY     | 0
TYPE             | 5
ERROR_CODE       | 0
NUM_ROWS         | 0
-[ RECORD 36 ]---
START_TIME       | 2010-10-14 12:25:02.142114
EPOCH_ID         | 434
NODE_NAME        | v_vmartdb_node01
USER_NAME        | dbadmin
SESSION_ID       | raster-s1-17172:0xa57
QUERY_IDENTIFIER |
QUERY            | GRANT SELECT ON sys_dba.QUERY_REPOSITORY TO PUBLIC;
SEARCH_PATH      | "$user", public, v_catalog, v_monitor, v_internal
DURATION         | 6634
SCHEMA_NAME      |
TABLE_NAME       |
PROJECTIONS      |
EXTRA_MEMORY     | 0
TYPE             | 5
ERROR_CODE       | 0
NUM_ROWS         | 0
-[ RECORD 37 ]---
START_TIME       | 2010-10-14 12:25:02.155941
EPOCH_ID         | 434
NODE_NAME        | v_vmartdb_node01
USER_NAME        | dbadmin
SESSION_ID       | raster-s1-17172:0xa58
QUERY_IDENTIFIER | query_repo-4
QUERY            | COPY sys_dba.query_repo (NODE_NAME,
SESSION_ID,QUERY_IDENTIFIER,QUERY,SEARCH_PATH,
DURATION,EPOCH_ID,START_TIME,PROJECTIONS,EXTRA_MEMORY,SCHEMA_NAME,TABLE_NAME,TYPE,
ERROR_CODE,USER_NAME,NUM_ROWS) FROM STDIN DELIMITER '|' DIRECT;
SEARCH_PATH      | "$user", public, v_catalog, v_monitor, v_internal
DURATION         | 125483
SCHEMA_NAME      | sys_dba
TABLE_NAME       | query_repo
PROJECTIONS      | sys_dba.query_repo_node01 sys_dba.query_repo_node02 sys_dba.query_repo_node03
                    sys_dba.query_repo_node04
EXTRA_MEMORY     | 0
TYPE             | 5

```

```

ERROR_CODE      | 0
NUM_ROWS        | 25
-[ RECORD 38 ]--+-+-----
...

```

See Also

Query Repository Parameters in the Administrator's Guide

Database Management Functions and QUERY_PROFILES in the SQL Reference Manual

Configuring Query Repository

Only the superuser can configure the query repository, which consists of:

- Configuring how and when data is saved to the query repository
- Configuring how and when data is cleared from the query repository
- Optionally setting the memory limit for the query repository

Configuring how and when data is saved to the query repository

The following parameters control how data is saved to the query repository:

Configuration Parameter	Default Value	Description
SaveQueryRepoInterval	300	Determines how frequently, in seconds, query data is saved to the query repository. The default is five minutes (300 seconds). Note: This parameter cannot be set to less than 300 seconds.
QueryRepoRetentionTime	7	Determines the maximum number of days of data that is saved in the query repository. This parameter also determines how much data, in days, Vertica retains in the query repository when it is cleared. The default value of this parameter is seven (7) days.
QueryRepoMemoryLimit	67108864	Determines the maximum memory available for the query repository and for storing database profiling data. The default is 64MB. When the cache becomes full, query repository data is dumped to a table. Profiling data is cleared from the cache. See Profiling Database Performance (page 11).

The following example determines that data is saved to the query repository every ten minutes (600 seconds).

```
=> SELECT SET_CONFIG_PARAMETER('SaveQueryRepoInterval', '600');
```

Note: Use the `SAVE_QUERY_REPOSITORY` function to trigger Vertica to save query data to the query repository immediately. Vertica saves data based on the established configuration parameters. For example, it uses the value of the `QueryRepoRetentionTime` parameter to determine the maximum number of days of queries to save. The `SAVE_QUERY_REPOSITORY` function resets the clock for the `SaveQueryRepoInterval` back to zero (0).

Configuring how and when data is cleared from the query repository

The following parameters control how data is cleared from the query repository:

Configuration Parameter	Default Value	Description
CleanQueryRepoInterval	86400	Determines how frequently, in seconds, query data is cleared from the query repository. The default is one day (86,400 seconds). This means that once daily any data that exceeds the query retention policy is cleared from the query repository.
QueryRepoRetentionTime	100	Determines the maximum number of days of query data to retain in the query repository each time data is saved. This parameter also determines how much data, in days, Vertica retains in the query repository when it is cleared. The default value of this parameter is seven (7) days.

The following example sets the QUERYREPORETENTIONTIME to 50 days:

```
=> SELECT SET_CONFIG_PARAMETER('QueryRepoRetentionTime', '50');
```

See Also

Configuration Parameters in the Administrator's Guide

Managing And Viewing Query Repository

Access to the query repository is strictly controlled. Full query repository data is available to the superuser only. Additionally, only the superuser has the ability to manage or monitor the query repository via:

- Configuration parameters
See **Configuring Query Repository** (page 23) and Configuration Parameters.
- Configuration functions
See CLEAR_QUERY_REPOSITORY and SAVE_QUERY_REPOSITORY.
- SELECT, DROP, and CREATE statements
See the SQL Reference Manual.

Users are restricted to viewing information about their own queries, which they access through the SYS_DBA.QUERY_REPO view. See Using Views in the Administrator's Guide.

Clearing the query repository immediately

The frequency with which the query repository is cleared is based on the CleanQueryRepoInterval parameter. You can also empty the query repository immediately.

To manually clear the query repository:

- 1 Note the value of the QueryRepoRetentionTime parameter.
- 2 Set the QueryRepoRetentionTime parameter to zero (0).
=> SELECT SET_CONFIG_PARAMETER('QueryRepoRetentionTime', '0');

- 3 Use the `CLEAR_QUERY_REPOSITORY` function to empty the query repository.
=> `SELECT CLEAR_QUERY_REPOSITORY();`
- 4 Set the `QueryRepoRetentionTime` parameter back to its original value before you changed it to zero. (The default value is 7.)
=> `SELECT SET_CONFIG_PARAMETER('QueryRepoRetentionTime','100');`

Using Diagnostic Tools

Vertica provides several diagnostic tools. In this section, you'll learn how to identify which version of Vertica you are running, use the Diagnostics utility, and export a catalog.

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 database logs along with other useful information into a single .zip file.

Syntax

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

Parameters

<i>command</i>	-h --help	Shows a help message and exits.
	-l --listdbs	Lists the running databases.
	-s --local_diags	Gathers diagnostics for local machine only.
	-d <i>dbname</i> --database <i>dbname</i>	Retrieves information about the database <i>dbname</i> .
	-o <i>OUTPUT_DIR</i> , --output_dir <i>OUTPUT_DIR</i>	Redirects output to somewhere other than /opt/vertica/log.
	-n <i>HOST_LIST</i> , --hosts <i>HOST LIST</i>	Gathers diagnostics for these hosts only.
	-z [all <i>n</i>]	Takes <i>all</i> or any positive integer value specifying how many <i>vertica.log-*.gz</i> files to include in diagnostics. The default is 1.

Notes

- In most cases, running `/opt/vertica/bin/diagnostics` without arguments is sufficient; the command reports on all databases and all nodes.
- The diagnostics utility writes the results into a file named `/opt/vertica/log/VerticaDiagnostics.date.zip` where *date* is replaced with a numeric date representation so that multiple calls to diagnostics result in different files being created.
- If you are having trouble with an installation, run the diagnostics utility as root or sudo. See **Running Diagnostics Utility for failed Installation** (page 43) for more information. For other situations, run the diagnostics utility as the database administrator.

- To include all *.gz files in diagnostics, use the `all` parameter, otherwise specify a number. The gzipped log files are included in the order in which they were last modified, with the most recent first.

The following command, for example, includes all `vertica*.gz` file in the diagnostics:

```
$ /opt/vertica/bin/diagnostics -d PROD01 -z all
```

The following command includes only the last three .gz files in the diagnostics:

```
$ /opt/vertica/bin/diagnostics -d PROD01 -z 3
```

- The diagnostics utility uses a `diagnostics-<username>.log` file instead of the `adminTools-<username>.log` file for logging to allow for improved readability of log files collected during diagnostics.

Exporting a Catalog

When you export a catalog you can quickly move a catalog to another cluster. Exporting a catalog transfers schemas, tables, constraints, projections, and views. System tables are not exported.

Exporting catalogs can also be useful for support purposes.

See the `EXPORT_CATALOG` function in the SQL Reference Manual for details.

Failure Recovery

Failure 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 recovers nodes by querying replicas of the data stored on other nodes. For example, a hardware failure could cause a node to lose database objects or to miss changes made to the database (INSERTs, UPDATEs, and so on) while offline. When the node comes back online, it recovers lost objects and catches up with changes by querying the other nodes.

Vertica uses the concept of K-Safety for failure recovery. The K value represents the maximum number of nodes in a database that can fail and recover with no loss of data. In Vertica, the value of K can be zero (0), one (1), or two (2). The Physical Schema design must meet certain requirements. To create designs that are K-Safe, Vertica recommends using the Database Designer.

The following table illustrates the number of nodes that can be down when the value of K is one (1) or two (2).

K Value	Nodes	Nodes DOWN	State of Database
1	3 or more	0	Safe.
		1	
		2	Unsafe. Automatic shutdown.
2	5 or more	0	Safe.
		1	
		2	
		3	Unsafe. Automatic shutdown.

Note: You can monitor the cluster state through the **View Database Cluster** state menu option.

Recovery Scenarios

Recovery comes into play when a node or the database is started. Depending upon how the node or database was shut down, there are three possibilities for a K-Safe database:

- **Recovery after failure of up to K Nodes:** This means that either one node has failed, for a database with a K value of one (1), or as many as two nodes have failed, for a database with a K value of two (2). Even in this state, the surviving 'UP' nodes of the database enable it to be fully operational and available for commands. The failed nodes can be restarted through the Administration Tools using the Restart Vertica on host option. The nodes being restarted show a status of 'RECOVERING' while they rebuild some of the data from the remaining nodes and, once finished, transition to an UP status. Transactions can continue to commit during the recovery process, except for a short period at the end of the recovery.

- **Recovery after a Clean Shutdown:** The database had been shut down cleanly via the Administration Tools Stop Database option. In this case, the database should be restarted using the Start Database option. Upon restart all nodes that were 'UP' at the time of shutdown immediately transition to 'UP'. It is possible that at the time of shutdown, the database could have had at most K(1) or K(2) nodes down. These nodes go through the 'RECOVERING' state as described in 'Recovery after failure of up to K nodes' case above.
- **Recovery after an Unclean Shutdown (Manual Recovery):** The database was not shut down cleanly, which means that the database lost more than K nodes and possibly did not write all the data from the WOS to disk. This could happen due to an various reasons; for example:
 - An unexpected event, such as a power failure that causes all nodes to reboot
 - Vertica processes on the nodes exited due to a software or hardware failure
 - Additional nodes failed during recovery, violating k-safety

When the database is started through the Administration Tools Start Database option, recovery determines that a normal startup is not possible. It goes on to determine a point in time in which the data was consistent on all nodes. This is called the Last Good Epoch. As part of Start Database processing, the administrator is prompted to accept recovery with the suggested epoch. If accepted, the database recovers and any data changes made after the Last Good Epoch are lost. If not accepted, startup is aborted and the database is not started on any of the nodes.

Instead of accepting the given epoch, the administrator can instead choose to recover from a backup or select an epoch for an even earlier point using the Roll Back Database to Last Good Epoch option in the Administration Tools Advanced Menu. This is useful in special situations, for example if the failure occurs during a batch of loads, for which it is easier to go back to the beginning of the batch, rather than starting in the middle, even though some of the work must be repeated. In most scenarios, it is sufficient and recommended to accept the given epoch.

Notes

- In Vertica 4.1, the default for the `HistoryRetentionTime` configuration parameter changed to 0, which means that Vertica only keeps historical data when nodes are down. This default setting effectively prevents the use of the Administration Tools 'Roll Back Database to Last Good Epoch' option because the AHM remains close to the current epoch and a rollback is not permitted to an epoch prior to the AHM. If you rely on the Roll Back option to remove recently loaded data, consider setting a day-wide window for removing loaded data; for example:

```
=> SELECT SET_CONFIG_PARAMETER ('HistoryRetentionTime', '86400');
```


See [Epoch Management Parameters](#) in the Administrator's Guide.
- Starting in 4.0, manual recovery is possible even if up to K nodes are out of commission; for example, physically removed for repair or not reachable at the time of recovery. Once the nodes are back in commission, they recover and rejoin the cluster, as described in the "Recovery after failure of up to K nodes" section above.
- **IMPORTANT:** When a node is down, it can take a full minute or more for the Vertica processes to time out during its attempt to form a cluster when manual recovery is needed. Wait approximately one minute until the system returns the manual recovery prompt. Do not press CTRL-C during database startup.

For information on troubleshooting recovery problems, refer to the following sections in the Troubleshooting Guide:

- **Startup Problems** (page 35)
- **Shutdown Problems** (page 31)

See Also

High Availability and Recovery in the Concepts Guide.

Shutdown Problems

A Database shuts down when one of the following events occurs:

- The administrator uses the Stop Database command.
- The cluster becomes unsafe and the database shuts down to prevent data loss.


If a problem prevents the database from shutting down, the Administration Tools display a dialog containing the following error message:

```
Database ... did not appear to stop...
```

The message is followed by a description of the problem. This section describes some of the known problems that can occur when stopping a database.

... large moveout is in progress

If there is a Tuple Mover operation in progress, the Administration Tools displays a message similar to the one shown below:



This particular message indicates that the Tuple Mover needs more time to complete a moveout operation, which is an internal session. (See Managing Sessions in the Administrator's Guide for more information.)

Note: 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 If you cannot wait any longer, disconnect and select Advanced > Stop Vertica on Host from the Administration Tools.
- 2 If Stop Vertica on Node fails, select Advanced > Killing a Vertica Process on Host. This forces the cluster to go through recovery at startup.

... users are connected

Error

If users are connected during shutdown operations, the Administration Tools displays a message similar to the following:

```
Database Stock_Schema did not appear to stop in the allotted time.
NOTICE: Cannot shut down while users are connected
        shutdown
```

```
-----
Shutdown: aborting shutdown
(1 row)
```

If you need to force a database shutdown, use the 'Stop Vertica on Node' command in the Advanced menu, selecting the appropriate nodes to stop.

Description

The message indicates that there are active user connections (sessions). See [Managing Sessions](#) in the Administrator's Guide for more information.

Resolution

The following examples were taken from a different database.

- 1 To see which users are connected, connect to the database and query the `SESSIONS` system table described in the SQL Reference Manual. For example:

```
=> \pset expanded
Expanded display is on.
=> SELECT * FROM SESSIONS;

-[ RECORD 1 ]
node_name      | site01
user_name      | dbadmin
client_hostname | 127.0.0.1:57141
login_timestamp | 2009-06-07 14:41:26
session_id     | rhel4-1-30361:0xd7e3e:994462853
transaction_start | 2009-06-07 14:48:54
transaction_id  | 45035996273741092
transaction_description | user dbadmin (select * from session;)
statement_start | 2009-06-07 14:53:31
statement_id    | 0
last_statement_duration | 1
current_statement | select * from sessions;
ssl_state       | None
authentication_method | Trust
-[ RECORD 2 ]
node_name      | site01
user_name      | dbadmin
client_hostname | 127.0.0.1:57142
login_timestamp | 2009-06-07 14:52:55
session_id     | rhel4-1-30361:0xd83ac:1017578618
transaction_start | 2009-06-07 14:53:26
transaction_id  | 45035996273741096
transaction_description | user dbadmin (COPY ClickStream_Fact FROM
'/data/clickstream/1g/ClickStream_Fact.tbl' DELIMITER '|' NULL '\n' DIRECT;)
statement_start | 2009-06-07 14:53:26
```

```

statement_id          | 17179869528
last_statement_duration | 0
current_statement     | COPY ClickStream_Fact FROM '/data/clickstream/1g/ClickStream_Fact.tbl'
DELIMITER '|' NULL '\\n' DIRECT;
ssl_state             | None
authentication_method | Trust

```

The current statement column of Record 1 shows that session is the one you are using to query the system table. Record 2 shows the session that must end before the database can be shut down.

- 2 If a statement is running in a session, that session must be closed. Use the function `CLOSE_SESSION` or `CLOSE_ALL_SESSIONS` described in the SQL Reference Manual.

Note: `CLOSE_ALL_SESSIONS` is the more common command because it forcefully disconnects all user sessions.

```
=> SELECT * FROM SESSIONS;
```

```

-[ RECORD 1 ]
node_name          | site01
user_name          | dbadmin
client_hostname    | 127.0.0.1:57141
client_pid         | 17838
login_timestamp    | 2009-06-07 14:41:26
session_id         | rhe14-1-30361:0xd7e3e:994462853
client_label       |
transaction_start  | 2009-06-07 14:48:54
transaction_id     | 45035996273741092
transaction_description | user dbadmin (select * from sessions;)
statement_start    | 2009-06-07 14:53:31
statement_id       | 0
last_statement_duration_us | 1
current_statement  | select * from sessions;
ssl_state          | None
authentication_method | Trust
-[ RECORD 2 ]
node_name          | site01
user_name          | dbadmin
client_hostname    | 127.0.0.1:57142
client_pid         | 17839
login_timestamp    | 2009-06-07 14:52:55
session_id         | rhe14-1-30361:0xd83ac:1017578618
client_label       |
transaction_start  | 2009-06-07 14:53:26
transaction_id     | 45035996273741096
transaction_description | user dbadmin (COPY ClickStream_Fact FROM
'/data/clickstream/1g/ClickStream_Fact.tbl'
DELIMITER '|' NULL '\\n' DIRECT;)
statement_start    | 2009-06-07 14:53:26
statement_id       | 17179869528
last_statement_duration_us | 0
current_statement  | COPY ClickStream_Fact FROM
'/data/clickstream/1g/ClickStream_Fact.tbl'
DELIMITER '|' NULL '\\n' DIRECT;
ssl_state          | None
authentication_method | Trust

```

```
=> SELECT CLOSE_SESSION('rhe14-1-30361:0xd83ac:1017578618');
```

```

-[ RECORD 1 ]
close_session | Session close command sent. Check sessions for progress.
=> SELECT * FROM SESSIONS;
```

```

-[ RECORD 1 ]
node_name          | site01

```

```
user_name          | dbadmin
client_hostname   | 127.0.0.1:57141
client_pid        | 17838
login_timestamp   | 2009-06-07 14:41:26
session_id        | rhel4-1-30361:0xd7e3e:994462853
client_label      |
transaction_start | 2009-06-07 14:48:54
transaction_id    | 45035996273741092
transaction_description | user dbadmin (select * from sessions;)
statement_start   | 2009-06-07 14:54:11
statement_id      | 0
last_statement_duration_us | 98
current_statement | select * from sessions;
ssl_state         | None
authentication_method | Trust
```

3 Query the SESSIONS table again. For example, two columns have changed:

- stmtid is now 0, indicating that no statement is in progress.
- stmt_duration now indicates how long the statement ran in milliseconds before being interrupted.

The SELECT statements that call these functions return when the interrupt or close message has been delivered to all nodes, not after the interrupt or close has completed.

4 Query the SESSIONS table again. When the session no longer appears in the SESSION table, disconnect and run the Stop Database command.

Controlling Sessions

The database administrator must be able to disallow new incoming connections in order to shut down the database. On a busy system, database shutdown is prevented if new sessions connect after the CLOSE_SESSION or CLOSE_ALL_SESSIONS() command is invoked — and before the database actually shuts down.

One option is for the administrator to issue the SHUTDOWN('true') command, which forces the database to shut down and disallow new connections. See SHUTDOWN in the SQL Reference Manual.

Another option is to modify the MaxClientSessions parameter from its original value to 0, in order to prevent new non-dbadmin users from connecting to the database.

1 Determine the original value for the MaxClientSessions parameter by querying the V_MONITOR.CONFIGURATIONS_PARAMETERS system table:

```
=> SELECT CURRENT_VALUE FROM CONFIGURATION_PARAMETERS WHERE
      parameter_name='MaxClientSessions';
CURRENT_VALUE
-----
50
(1 row)
```

2 Set the MaxClientSessions parameter to 0 to prevent new non-dbadmin connections:

```
=> SELECT SET_CONFIG_PARAMETER('MaxClientSessions', 0);
```

Note: The previous command allows up to five administrators to log in.

3 Issue the CLOSE_ALL_SESSIONS() command to remove existing sessions:

```
=> SELECT CLOSE_ALL_SESSIONS();
```


4 Query the `SESSIONS` table:

```
=> SELECT * FROM SESSIONS;
```

When the session no longer appears in the `SESSIONS` table, disconnect and run the Stop Database command.

5 Restart the database.

6 Restore the `MaxClientSessions` parameter to its original value:

```
=> SELECT SET_CONFIG_PARAMETER('MaxClientSessions', 50);
```

See Also

`CLOSE_SESSION`, `CLOSE_ALL_SESSIONS`, `CONFIGURATION_PARAMETERS`, and `SESSIONS` in the SQL Reference Manual

Managing Sessions and Configuration Parameters in the Administrator's Guide

No running statement, that session is idle

Error

```
No running statement, that session is idle
```

Description

The `INTERRUPT_STATEMENT` function failed. The session exists but is not running a statement and the session ID can't be found.

Resolution

Not required.

Startup Problems

This section describes some of the known problems that can occur when starting a database. Startup fails on one or more nodes either due to communication problems between nodes or when recovery fails to recover the data on the node(s) for any reason.

Note: 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

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

Description

This message typically indicates an abnormal shutdown of one or more nodes. If any nodes in the database are UP, the database is fully operational. Note however that if you start load processing on the database the recovery node could take additional time to become operational

Resolution

None required

Error starting database, no nodes are up

Error

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

Description

An unknown problem is preventing the database from starting.

Resolution

Reboot the hosts and try to start the database. If unsuccessful, run diagnostics and contact **Technical Support** (on page 1).

Database startup successful, but it could be incomplete

Error

```
Database startup successful, but it may be incomplete.
```

Description

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

Resolution

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

Error

```
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
```

Description

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

Resolution

If this error persists, contact **Technical Support** (on page 1).

Database did not start cleanly on initiator node!

Error

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

Description

Configuration problems can cause this error.

Resolution

- 1 Check hostname resolution as described in Configure 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`
- 3 Verify that there is no firewall running.

TIMEOUT ERROR: Could not login with SSH

Error

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

Description

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

Resolution

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

Good epoch logs are available on all nodes

Error

```
Database startup failed. Good epoch logs are available on all nodes.  
WARNING: if you say 'yes.' changes made to the 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)?
```

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 allows the cluster to proceed normally using the data saved prior to the epoch.

No good epoch log available on node

Error

```
Database startup failed. No good epoch log available on node stock_multi_node_0.  
Please run diagnostics and contact Vertica
```

Description

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.

Resolution

- 1 Make sure that all nodes are powered on.
- 2 Start the database again.
- 3 Make sure that all nodes have Spread running. If necessary, restart Spread where it is not running and start the database.
- 4 On each node that did not start up, examine dbLog for the cause of the failure.
- 5 If the cause cannot be determined, it is likely that a node has no catalog version or epoch log from which to recover. Run diagnostic tests (see *Using Diagnostic Tools* (page 26)) and contact *Technical Support* (on page 1).

Nodes stuck in INITIALIZING state

Error

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

Description

This issue is known to happen due to network configuration problems, corrupted catalogs or disks, missing database directories or some other fatal problem in the database setup. Incorrect use of the Administration Tools Advanced Menu options could also lead to this condition.

Resolution

- 1 Open the Administration Tools.
- 2 Select Advanced > Stopping Vertica on Host to stop all nodes.
- 3 Go back to the Main Menu and click Start Database.
- 4 If the error persists, run diagnostics and contact *Technical Support* (on page 1)

Node does not recover because of lock timeouts

Error

Upon starting a node, it stays in RECOVERING state for a long time and eventually shuts down again. Examination of the *vertica.log* on the node reveals an error:

```
Locking failure: Timed out S locking Table.
```

Description

The final stage of recovering a node requires a S lock on the table. If you have a continuous stream of COPY commands in progress, recovery might not be able to get this lock even after multiple re-tries.

Resolution

If you see this situation, either momentarily stop the loads or pick a time when the cluster is not busy to restart the node and let recovery proceed.

Spread Problems

This section describes some of the known problems that can occur when using spread.

Spread is not running

Error

```
admintools View Cluster State shows "Could not connect to spread.
Spread is configured as part of database creation".
spread dead but pid file exists
```

Resolution

Verify that spread is not running and then restart the spread daemon.

- 1 Verify that spread is not running:


```
ps ax | grep spread
```
- 2 Examine */tmp/spread*.log* and */var/log/spreadd.log* for problems. Permission problems and syntax problems are identified in these log files.
- 3 Issue the *ifconfig* command to check the current IP addresses of the hosts and verify that those IP addresses are listed in */opt/vertica/config/vspread.conf*.
- 4 Check for Vertica processes that might be running, even though spread is down:


```
ps ax | grep vertica
```
- 5 Kill the Vertica process. Alternatively, use the Admintools Advanced Menu > **Kill Vertica Process on Host**.
- 6 Restart spread:
 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. Ensure the daemon is running:  
# ps ax | grep spread
```

Administration Tools shows node state as UNKNOWN

Error

The Administration Tools 'View Cluster State' shows one or more nodes with an UNKNOWN status

Description

Under certain conditions, Vertica nodes can go into the UNKNOWN state yet still be processing. In most cases, after some time, they return to UP status. However, if you see a persistent UNKNOWN state that does not resolve to an UP state after several minutes, follow the instructions in this section.

Resolution

Likely cause of this issue is a sub-optimally configured I/O subsystem leading to high contention that causes Vertica to be unresponsive to messages from the spread daemon. You might notice this problem occurs more readily when running several large join queries that spill to disk. Check for high I/O waits and other symptoms of I/O problems and if unable to resolve, contact **Technical Support** (on page 1) with I/O statistics and sar data.

See Also

Spread is not running (page 39) for restarting spread

Diagnosing spread problems

Spread Panic

Error

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

Description

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.

Resolution

Do not attempt to change the spread configuration. Contact **Technical Support** (on page 1).

Spread Dead but pid File Exists

Error

```
spread dead but pid file exists
```

Description

If spread ends abnormally, the pid and lock file are left behind.

Resolution

Restart spread.

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 Ensure the daemon is running:

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

To diagnose issues related to starting Spread, "status" option has been enhanced to provide guidance .

Example

The following example checks on the spread status.

```
$ sudo /etc/init.d/spreadd status
```

```
spread is stopped
```

If you are having trouble starting spread, check `/opt/vertica/config/vspread.conf` and spread logs in `/tmp/spread_*` and `/var/log/spreadd.log`.

```
$ sudo /etc/init.d/spreadd start
```

```
Starting spread daemon: spread (pid 24290) is running...
```

```
[ OK ]
```

```
$ sudo /etc/init.d/spreadd status
```

```
spread (pid 24290) is running...
```

Installation Problems

This section describes installation issues you might encounter.

Failed Dependencies

Error

```
ERROR: Failed dependencies:
```

Description

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

Resolution

Install the missing product. For example:

Python

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

Note: 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>).

Suse10 and libtermcap.so

On SUSE 10, the Vertica install could fail due to a failed dependency check for missing libtermcap.so:

```
ERROR: Failed dependencies:
libtermcap.so.2()(64bit) is needed by vertica-4.0.11-20100503180005.x86_64
```

To resolve this issue:

- 1 Download the `termcap*.rpm` from the **openSUSE Web site** <http://software.opensuse.org>.

- 2 Install the rpm on the SUSE cluster as root or sudo:

```
# sudo rpm -ivh termcap-2.0.8-910.x86_64.rpm
```

Vertica Web Monitoring (Ganglia)

The Vertica Monitoring (Ganglia) rpm install could fail if it encounters a missing dependency. These dependencies come pre-installed on most Linux distributions; however, some could be missing on older Linux distributions.

The following is an example of a failed dependency:

```
error: Failed dependencies:
  ganglia-gmetad is needed by vertica-ganglia-1.1.0-1.noarch
  ganglia-gmond is needed by vertica-ganglia-1.1.0-1.noarch
  php is needed by vertica-ganglia-1.1.0-1.noarch
```


php-gd is needed by vertica-ganglia-1.1.0-1.noarch

See Required Packages in the Administrator's Guide for guidelines on resolving missing dependencies.

SSH Already Configured

Error

SSH already configured

If root does not have password-less you are prompted three times for each host

Description

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

Resolution

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.

Running Diagnostics Utility for Failed Installation

Description

To troubleshoot an installation problem, run the diagnostics utility as root with the `--hosts` option.

Example

```
/opt/vertica/sbin/diagnostics
--hosts node01,node02,node03
```

SSH logins must be enabled on all nodes in order for the diagnostics utility to run.

If the `install_vertica` script failed to enable SSH logins, then run the ***diagnostics utility*** (page 27) as root or sudo individually on each node.

Diagnostics

Contact ***Technical Support*** (on page 1) and provide the diagnostics .zip files.

Passwordless SSH Access to Other Sites Failed

Error

Passwordless SSH access to other sites ... FAILED

Description

For database administration to function properly, Vertica requires that passwordless authentication for SSH be configured for the Vertica DBA user. The Vertica install script attempts to set this up automatically for the DBA user during installation. However, this process could fail if:

- The Vertica DBA user already exists and has conflicting SSH keys.
- SSH is not configured correctly to allow passwordless authentication.
- Permissions changed on the home directory for dbadmin (`/home/dbadmin`) to allow others to modify (e.g., `755`); thus, `sshd` refuses to allow passwordless authentication using public (ssh) keys (although identification using normal password works).

Resolutions

- Verify that `/etc/ssh/sshd_config` is configured properly to allow passwordless authentication using public keys.
- Verify that the `AuthorizedKeysFile` setting in `/etc/ssh/sshd_config` is either set to `"authorized_keys2"` or commented out. After changing the value of `AuthorizedKeysFile`, you will need to restart `sshd`:

```
/etc/init.d/sshd restart
```
- Reset permissions on the home directory so that it is read-write-execute by user dbadmin only. For example:

```
# chmod 700 ~
```

See Also

Enable Secure Shell (SSH) Logins and Provide Root SSH Access to the Cluster in the Installation Guide

Misleading Message Due to IP/hostname Mismatch

Error

```
Installation Tool EOF ERROR: Could not login with SSH. Here is what SSH said:  
ssh: connect to host ellis port 22: Connection refused  
User root cannot login to ellis  
Installation failed.
```

Description

Vertica requires IP addresses for communication among cluster hosts to be statically assigned.

If you happened to be using DHCP and the IP addresses in your `/etc/hosts` do not match the actual IP address assigned by DHCP for the hosts, you could get an error as described above.

Resolution

Do not use DHCP to assign IP addresses for network interfaces used to communicate between Vertica hosts.

If you use separate public and private interfaces for external and internal network communications to the host, you can directly specify the private IP addresses as parameters to `install_vertica` eliminating the dependency on hostnames defined in `/etc/hosts`.

Host Key Verification Failure

Error

```
EOF ERROR: Could not login with SSH. Here is what SSH said:
WARNING: REMOTE HOST IDENTIFICATION HAS CHANGED!
...
Host verification failed.
```

Description

After a host is lost and replaced in the Verticacluster, the host key will change and the `known_hosts` files on the surviving nodes have an old value. Therefore, when you use the Administration Tools to restart the newly-added node, the system returns the above warning.

Resolution

Remove the old host key from `~dbadmin/.ssh/known_hosts`.

Permission denied even after providing correct root password

Error

During the installation, you were prompted three times for the root password, but even after password was supplied, access was denied. You see the following error message:

```
ERROR: Permission denied. Incorrect password. Please check that your password is
correct and that the user has permission to login via ssh
```

Resolution

The `PermitRootLogin` setting in `/etc/ssh/sshd_config` must be set to "yes" or the `PermitRootLogin` setting must be commented out.

After changing the value of `PermitRootLogin`, you will need to restart `sshd`:

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

Installation runs as sudoer when run using sudo root

Error

When logged in as a normal user (for example, `myuser`), if installation was run by issuing `sudo su root` to change to root user, it appears to still run as `myuser`.

The `admintools.log` shows that SSH is being done as `myuser` rather than root:

```
Sep 25 22:50:05 [adminExec.__init__] Unique Ports: False
Sep 25 22:50:05 [SSH.login data dev2]
Sep 25 22:50:05 ssh -l myuser dev2
```

```
Sep 25 22:50:05 [SSH.login ssh -l myuser dev2 ]
Sep 25 22:50:06 Password attempt #2 2
Sep 25 22:50:06 Password attempt #4 0
Sep 25 22:50:15 Password attempt #3 -1
Sep 25 22:50:15 myuser@dev2: sudo [ -f /etc/SuSE-release ]
Sep 25 22:50:16 myuser@sdev2: ['1', ['']]
Sep 25 22:50:16 Error code 1 ['']
...
```

Resolution

`sudo su root`, `sudo -i`, and `sudo /opt/vertica/sbin/install_vertica` all end up having the `SUDO_UID` values set to the sudoer user, such as `myuser` in the above example. The installation script uses the presence of this environment variable to determine if it should run in 'sudo' mode or standard "root" mode. In these instances, it thinks it is running as the `SUDO_UID` user and assumes that the root password is not available; thus it runs all commands under the sudo user.

When switching to root or another user using `sudo`, use the `-` (hyphen), which completely replaces the environment and `SUDO_UID` is not set.

Managing Large Grouped ROS Containers

If you upgrade your database from Vertica 4.0.x to 4.0 Patchset 1, check for large grouped ROS containers, especially if you used `COPY AUTO` or `COPY` [no keyword] for bulk loads. Note that this issue does not affect a 3.5.x database.

Small Partitioned Tables

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

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

Important: For large unpartitioned tables, contact Vertica *Technical Support* (on page 1).

Large Partitioned Tables

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. Contact **Technical Support** (on page 1) for guidance.

See Also

DO_TM_TASK and MERGE_PARTITIONS in the SQL Reference Manual

Backup and Restore Problems

This section describes backup and restore issues you might encounter.

Errors found in SAL or Catalog

Error

Errors found in SAL or Catalog, check log for details

Description

A catalog snapshot file that you attempted to use to restore a database is corrupt. This can occur for a number of reasons, which include:

- The snapshot file is bad.
- The correct data files were not copied to the correct locations.

Resolution

See the `bootstrap-catalog.log` for details.

Cannot bootstrap catalog on top of existing catalog

Error

Cannot bootstrap catalog on top of existing catalog

Description

A catalog directory with catalog information exists in the location specified for the new catalog (`dbdir/catalog`).

Resolution

Verify that you want to replace the original catalog. To force a replacement, use the `-O` flag when you initialize the backup catalog:

```
bootstrap-catalog -O -D <dbdir> -F <snapshotfile>
```

Using the `-O` flag renames the existing catalog directory to a temporary name (`Catalog-old-<pid>`), initializes the new catalog, and then removes the old catalog directory. If the new catalog fails to bootstrap, the old catalog directory is not removed, thus allowing you to restore the database node to its original state.

Restore fails if files present in catalog directory

Description

The restore (`restore.sh`) script fails to restore the database if core files are present in Vertica database directory.

Resolution

Certain files names are allowed to be present in the top-level catalog directory for a successful bootstrap/restore. Any other files cause the bootstrap/restore to fail.

Be sure no core files are present with the exception of the following files, which are allowed and/or needed:

- Anything starting with vertica.log
- Epoch.log
- vertica.pid
- not-yet-initialized
- global
- ErrorReport.txt
- SAL
- tmp
- CopyErrorLogs
- Snapshots
- Anything starting with Catalog-old-
- bootstrap-catalog.log
- Anything ending in .conf

Configuration Problems

This section describes configuration issues you might encounter.

Administration Tools Errors

Description

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.

Resolution

Change the shell prompt to use standard bash type characters.

Cannot Allocate Memory

Error

```
Cannot allocate memory
```

Description

There are four known causes for this:

- 1 max_map_count set too low. See Increase the max_map_count Parameter.
- 2 max memory size (ulimit -m) set too low.
- 3 virtual memory (ulimit -v) set too low.
- 4 /proc/sys/vm/overcommit_memory set to 2 and CommitLimit (see /proc/meminfo) set too low.

Resolution

For Vertica to function properly, max memory size and virtual memory should be unlimited and /proc/sys/vm/overcommit_memory should be set to 0.

See your operating system documentation for how to set these values. You need to set them such that the value you specify is used even after reboot.

Cannot find IP Address for Host...

Error

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

Description

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

Resolution

- 1 See Configure 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

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

Description

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

Resolution

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.

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>

Could Not Log in with SSH

Error

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

Description

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

Resolution

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

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

Resolution

Make sure that PostgreSQL is not installed on any host.

Create Database Fails

Error

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

Description

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

Resolution

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.

Denied. Not a super user. Cannot create a schema

Error

```
Denied. XX is not a superuser, so cannot create a schema for YY
```

Description

A user other than the superuser attempted to assign ownership of a schema to another user. Only the superuser can assign ownership of a schema to another user.

Resolution

The superuser assigns ownership of the schema to the user.

Error Detected During Node Verification

Error

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

Resolution

Check Configure Hostname Resolution in the Installation Guide.

Insufficient Privilege USAGE on SCHEMA

Error

Insufficient privilege: USAGE on SCHEMA XX not granted for current user

Description

A user who has not been granted the USAGE privilege for the schema attempted to access a table within the schema. To access schema objects, the user must be granted USAGE privilege on the schema in addition to privileges for specific objects within the schema.

Resolution

Grant the user USAGE on the schema.

Not Enough Open File Handles Allowed...

Description and Resolution

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>
```

where `<value>` = 1 file per 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 1 file per 1 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 Vertica:

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

```
* soft NOFILE <value>
* hard NOFILE <value>
```

where `<value>` = 1 file per 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 once again.

- 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'.
```

See Also

Verify that Kernel and User Parameters Were Set

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

Using sudo NOFILE Parameters Not Saved in limits.conf

Issue and Resolution

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

Location	Setting	Value	Comment
/etc/security/limits.conf	NOFILE	65536	Sets the maximum number of open files for the user.
/etc/security/limits.conf	NPROC	Total RAM on the node in MB (e.g., on 1GB machine, set NPROC to 1000)	Controls the maximum number of threads and processes for each user.
/etc/sysctl.conf	fs.file-max	65536	Sets the maximum number of file handles that the Linux kernel will allocate.
/etc/sysctl.conf	vm.max_map_count	Total RAM on the node in KB/16	The maximum number of memory map areas a process might have.
/etc/rc.local	blockdev --setra	2048	Sets the readahead parameter.
/etc/sysctl.conf	vm.min_free_kbytes	4096	Changes the page reclaim thresholds. When this number is increased the system starts reclaiming memory earlier, when its lowered it starts reclaiming memory later.
/etc/security/limits.conf	fsize	unlimited	Sets the maximum file size to the OS limit

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 Install](#).

- Increase the Maximum Number of Files Open - use this procedure for any settings in `/etc/security/limits.conf` settings.

- Increase the `max_map_count` Parameter - use this procedure for any settings in `/etc/sysctl.conf`.

Unable to drop table. Must be Owner of Relation

Error

```
Must be owner of relation XX
```

Description

A user attempted to drop a table without the appropriate privilege. To drop a table, the user must be granted ALL privileges on the schema and must own the table.

Resolution

The table owner with ALL privileges on the schema drops the table.

I/O Error: Permission Denied

Error

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

Description

The configuration directory is writable only by root.

Resolution

- 1 Log in as root:

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

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

Permission Denied for Relation

Error

```
Permission denied for relation XX
```

Description

A user attempted to perform a SELECT, UPDATE, or DELETE operation on a table without the appropriate privilege. The user must be granted appropriate privileges on the schema and table before the user can perform the operation.

Resolution

Grant the user appropriate privileges on the schema and table.

One or More Nodes Did Not Open a Data Connection to this Node

Error

```
ROLLBACK: One or more nodes did not open a data connection to this node.
```

Description

This could indicate a network configuration problem. Check that the private interfaces used for communication among the cluster hosts reside in the same subnet and are returned first by host address lookup. If using */etc/hosts* for host name resolution, place the private interfaces used by the database above any other addresses with the same host names in */etc/hosts*. You can verify the host names by logging in to each node and pinging each of the other nodes. Verify that the host name reported back by ping matches the node's host name and is the same as any host name you used for your configuration when you installed Vertica.

Resolution

All nodes in the database cluster must be in complete agreement about host names and addresses.

Spread Dead but pid File Exists

Error

```
spread dead but pid file exists
```

Description

If spread ends abnormally, the pid and lock file are left behind.

Resolution

Restart spread.

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 Ensure the daemon is running:

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


To diagnose issues related to starting Spread, "status" option has been enhanced to provide guidance .

Example

The following example checks on the spread status.

```
$ sudo /etc/init.d/spreadd status
```

```
spread is stopped
```

If you are having trouble starting spread, check `/opt/vertica/config/vspread.conf` and spread logs in `/tmp/spread_*` and `/var/log/spreadd.log`.

```
$ sudo /etc/init.d/spreadd start
```

```
Starting spread daemon: spread (pid 24290) is running...
```

```
[ OK ]
```

```
$ sudo /etc/init.d/spreadd status
```

```
spread (pid 24290) is running...
```

Connection Problems

This section describes connection issues you might encounter.

Could Not Connect to Database

Error

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

Description

You probably have hostname resolution problems.

Resolution

Make sure you have followed the instructions in Configure Hostname Resolution in the Installation Guide.

Could Not Connect to Server: Connection Refused

Error

```
could not connect to database
connection refused
```

Description

The database is not up or not listening on the port you are connecting to, or there are hostname resolution problems on the network interface used to connect to the database by the client.

Resolution

Check View Cluster State to ensure that the database node you are connecting to is UP.

Make sure you have followed the Configure Hostname Resolution instructions in the Installation Guide.

New Session Rejected Due to Limit

Error

```
New session rejected due to limit, already 50 sessions active
```

Description

By default Vertica allows 50 client sessions and an additional 5 administrator sessions. The system returns a message if the database exceeds the limit.

Resolution

Modify the `MaxClientSessions` parameter. For example, to increase the number of `MaxClientSessions` to 100, issue the following series of commands at a vsql prompt:

- 1 Determine the original value for the `MaxClientSessions` parameter:

```
=> SELECT CURRENT_VALUE FROM CONFIGURATION_PARAMETERS WHERE
      parameter_name='MaxClientSessions';
      CURRENT_VALUE
      -----
      50
      (1 row)
```

- 2 Increase the `MaxClientSessions` parameter to 100.

```
=> SELECT SET_CONFIG_PARAMETER('MaxClientSessions', 100);
```

To set `MaxClientSessions` to the original value, issue the following command.

```
=> SELECT SET_CONFIG_PARAMETER('MaxClientSessions', 50);
```

See Also

... *users are connected* (page 32)

Managing Sessions and Configuration Parameters in the Administrator's Guide

`CLOSE_SESSION`, `CLOSE_ALL_SESSIONS`, `CONFIGURATION_PARAMETERS`, and `SESSIONS` in the SQL Reference Manual

Firewall Configuration

Vertica does not recommend using firewalls between nodes. If you are running your cluster on a private interface, a firewall is not needed. A firewall on the public side should only need port 5433 open.

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)

Note: On the private interface, in addition to those listed, all the dynamic ports, which are often between 32768 and 61000 but depend on the OS configuration, would need to be open. Also SSH (port 22). Again, Vertica does not recommend having firewall rules for the private (Vertica) networks.

Firewall inactivity timeout hit (Insufficient KEEPALIVE value)

Issue

The connection between the server and client is killed prematurely by a firewall timeout. This could happen when a long-running query is in progress but no data is being passed back to the client, and the firewall timeout is less than the TCP KEEPALIVE setting on the database server. (On some Linux distributions, the default is 2 hours or 7200 seconds.)

Resolution

Setting the KEEPALIVE setting to a value lower than the firewall timeout might help. In the following example, it is being set to 10 minutes (600 seconds).

```
echo 600 > /proc/sys/net/ipv4/tcp_keepalive_time
```

Database Designer Errors

This section describes Database Designer errors you might encounter.

No Such File or Directory ... super_projections.tmp.sql

Error

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

Resolution

Create the directory manually. For example:

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

Startup Problems

This section describes startup issues you might encounter.

KeyError: 'getpwuid(): uid not found: <uid>'

Error

```
KeyError: 'getpwuid(): uid not found: <uid>'
```

Description

A database that needed automatic system recovery was started through the Administration Tools, but timed out when the prompt to start the server at the last good epoch was ignored. This occurred because the idle time exceeded the value of the SSH daemon configuration file (`sshd_config`).

Resolution

Log out and then log back in to start a fresh session.

(LOSTCONTACT) Slow Database Startup when Node is Down

Error

No message.

Description

When a node is down, it can take a full minute or more for the Vertica processes to time out during its attempt to form a cluster when manual recovery is needed.

Resolution

Wait approximately one minute until the system returns the manual recovery prompt. Do not press CTRL-C during database startup.

Vertica Fails to Start (verticad and pam_limits)

Description

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 `verticad` init script would fail to start 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.

Resolution

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.

Storage Location Problems

This section describes storage location issues you might encounter.

Could Not Open Directory

Error

```
could not open directory  
<directory_path>: Not a directory
```

Description

You cannot add a location to a non-existent directory

Resolution

Verify the directory path where you want to create the storage location. If the directory doesn't exist, create it before attempting to create the storage location.

Could Not Create Directory

Error

```
could not create directory  
<directory_path>: Permission denied
```

Description

You do not have sufficient permission to create a directory.

Resolution

The person who creates the directory must have system privileges for creating directories.

Could Not Add Location (Permission)

Error

```
could not add location  
<directory_path>: Permission denied
```

Description

You cannot add a storage location because you do not have read/write permission on the parent directory.

Resolution

The person who creates the storage location must have read/write permission on the parent directory.

Could Not Add Location (Unavailable)

Error

```
could not add location  
<directory_path>: Directory not empty
```

Description

You cannot add a storage location to a directory that already contains files.

Resolution

Either create a new directory in which to create a storage location or move the existing files to another directory.

Usage Cannot be an Empty String

Error

```
usage cannot be an empty string
```

Description

When using the ALTER_LOCATION_USE function to alter the use of a storage location, you did not specify its new use.

Resolution

Be sure to specify the usage string.

The following example alters the storage location on node3 to store data only:

```
=> SELECT ALTER_LOCATION_USE ('/thirdVerticaStorageLocation/' , 'node3' ,  
'DATA');
```

Usage Cannot be Changed to Data

Error

```
ROLLBACK: <storage_location> cannot be changed to data. There will be no storage  
locations for temporary files.
```

Description

You attempted to change the usage for the only storage location that stores temporary files.

Resolution

At least one storage location must be available to store temporary files. Create a new location to store temporary files before you attempt to change the usage of this one. See ADD_LOCATION.

Usage Cannot be Changed to Temp

Error

ROLLBACK: <storage_location> cannot be changed to temp. There will be no storage locations for data files.

Description

You attempted to change the usage for the only storage location that stores data files.

Resolution

At least one storage location must be available to store data files. Create a new location to store data files before you attempt to change the usage of this one. See `ADD_LOCATION`.

Usage of Storage Location Cannot be Changed

Error

Usage of <storage_location> cannot be changed. It has been retired.

Description

You attempted to change the usage for a storage location that has been retired. Vertica prohibits this.

Throughput Should be Greater Than (>) Zero (0)

Error

```
throughput should be > 0
```

Description

When specifying the performance for the storage location, you specified a throughput of less than 1.

Resolution

The throughput of the storage location performance must be 1 or more. See `SET_LOCATION_PERFORMANCE`.

Latency Should be Greater Than (>) One (1)

Error

```
latency should be > 0
```

Description

When specifying the performance for the storage location, you specified a latency of less than 1.

Resolution

The latency of the storage location performance must be 1 or more. See `SET_LOCATION_PERFORMANCE`.

Cannot be Retired**Error**

```
ROLLBACK: <storage_location> cannot be retired.  
There will be no storage locations for temporary/data files.
```

Description

You attempted to retire the only storage location that stores temporary or data files.

Resolution

At least one storage location must be available to store temporary or data files. Create a new location to store temporary or data files before you attempt to retire this one. See `ADD_LOCATION`.

DROP_LOCATION for Data Locations is not Supported**Error**

```
drop_location for DATA locations is not supported.
```

Description

You attempted to drop a location that is being used to store data files. Vertica prohibits dropping locations for which the usage is set to `DATA` or `DATA,TEMP`.

Resolution

Change the usage of the location to `TEMP` and verify that no data files remain. If one or more data files remain, you can perform a manual moveout of the data files, wait for the ATM to merge out the data files automatically, or you can drop partitions. Deleting data files does not work.

Note: At least one storage location must be available to store data files.

Location Cannot be Dropped**Error**

```
Location cannot be dropped as it stores data files.
```

Description

You attempted to drop a location that contains data files. This occurs when the usage for a storage location used to be set to `DATA` or `DATA,TEMP` and the location still contains data files.

Resolution

If one or more data files remain, you can perform a manual mergeout of the data files, wait for the ATM to merge out the data files automatically, or you can drop partitions. Deleting data files does not work.

Execution Problems

This section contains error messages related to SQL execution issues.

Attempted to Create Too Many ROS Containers

Error

Attempted to create too many ROS containers for projection

Description

Heavy data load conditions can cause the Tuple Mover to fall behind in performing moveout or mergeout operations. The resulting large number of ROS containers can cause some requests to exhaust all available system resources. Vertica detects this problem and prevents load transactions until the Tuple Mover has time to catch up.

Resolution

You might need to adjust the Tuple Mover's configuration parameters to compensate for the load pattern or rate. See [Tuning the Tuple Mover in the Administrator's Guide](#) for details.

You can query the `TUPLE_MOVER_OPERATIONS` table to monitor mergeout activity. However, the Tuple Mover does not immediately start a mergeout when a projection reaches the limit of ROS containers, so you may not see a mergeout in progress when receiving this error.

If waiting for a mergeout does not resolve the error, the problem probably is related to insufficient RAM. In that case, contact [Technical Support](#) (on page 1). A good rule of thumb is that system RAM in MB divided by 6 times the number of columns in the largest table should be greater than 10. For example, for a 100 column table you would want at least 6GB of RAM ($6144\text{MB} / (6 * 100) = 10.24$) to handle continuous loads.

Cannot Drop Partition with Nodes Down

Error

ROLLBACK: Can't drop a partition with nodes down

Description

This error is obsolete. If a node fails during drop operations, `DROP_PARTITION` commits the dropped partition during recovery.

See Also

[DROP_PARTITION](#) in the SQL Reference Manual

[Partitioning Tables](#) in the Administrator's Guide

Case-sensitive System Table Queries Return 0 Records

Some system table data may be stored in mixed case. For instance, Vertica stores mixed-case identifier names as they were specified in the CREATE statement, even though the case is ignored when they are referenced in queries. See Identifiers. When these object names appear as data in the system tables, it is error prone to retrieve them with the equality (=) predicate because the case must match exactly to what is stored. It is much easier to use the case-insensitive operator ILIKE instead.

Example:

Given the following schema:

```
=> CREATE SCHEMA SS;
=> CREATE TABLE SS.TT (c1 int);
=> CREATE PROJECTION SS.TTP1 AS SELECT * FROM ss.tt UNSEGMENTED ALL NODES;
=> INSERT INTO ss.tt VALUES (1);
```

If you run a query using the = predicate, Vertica returns 0 rows:

```
=> SELECT table_schema, table_name FROM v_catalog.tables WHERE table_schema = 'ss';
table_schema | table_name
-----+-----
(0 rows)
```

Using the case-insensitive ILIKE predicate returns the expected results:

```
=> SELECT table_schema, table_name FROM v_catalog.tables WHERE table_schema ILIKE
'ss';
table_schema | table_name
-----+-----
SS           | TT
(1 row)
```

Column Definition Changed to NOT NULL

Error

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

Description

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

Resolution

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

COPY, DELETE, UPDATE, or INSERT Hangs When Server Idle

Resolution

Check for an idle session that is holding a lock, such as a SELECT statement without AT EPOCH LATEST. Review all of the sessions shown by

```
=> SELECT * FROM VT_SESSION
```

Consider using the READ COMMITTED transaction model, which gives the same results as always using AT EPOCH LATEST

Could Not Write To Filename

Error

```
ROLLBACK, ERRCODE_DISK_FULL
Could not write to filename: reason
where the reason is of the form:
Volume path has amount in bytes bytes free (amount in bytes unreserved).
Minimum free space is amount in bytes (type of limit)
where the type of limit is one of:
    Catalog
    Table Data
    Temporary Data
```

Description

The database has detected a low disk space condition.

Resolution

See Managing Disk Space in the Administrator's Guide

Ctrl-C Distorts Query Result Set Output

Description

If you interrupt an interactive query by typing Ctrl-c after the result set starts to display, your terminal settings could be 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]

Resolution

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

Deleting Records from Table with PK with Associated FK

Error

Deleting the records from a table with a primary key does not delete the records from a table with the associated foreign key

Description

If you delete records from a dimension table, any records in a fact table or an associated dimension table are not automatically deleted

Errors that Result in a Statement-Level Rollback

A statement-level rollback reverses just the effects made by a particular statement. Most errors caused by a statement result in a statement-level rollback to undo the effects of the erroneous statement. Vertica uses ERROR messages to indicate this type of error.

See Transactions.

Errors that Result in a Transaction-Level Rollback

DDL errors, systemic failures, deadlocks, and resource constraints result in transaction-level rollback, which discards all modifications made by a transaction. Vertica uses ROLLBACK messages to indicate this type of error.

See Transactions.

Execution Error on External Procedures

Error

```
INFO: Procedure reported:  
Could not execute procedure, errno=8  
ERROR: Procedure execution error: exit status=1
```

Description

When executing an external procedure from the command line, the client returns an error.

Resolution

- Make sure there are no extraneous spaces before the command `#!/bin/bash`.
- If you are using programs within the external procedure, such as `vsql`, make sure the program is fully pathed in the script—even if it's already in the search path; for example `/opt/vertica/bin/vsql`.

Function <function-name>(<data-type>, ...) does not exist

Error

```
ERROR: function convert(varchar, timestamptz, int) does not exist
```

Description

Vertica returns a syntax error if anything is entered that appears to have a type as an argument, rather than an expression as an argument.

Resolution

Use only expressions as arguments to SQL functions.

Function with specified name and parameters does not exist

Error

```
ROLLBACK: Function with specified name and parameters does not exist
```

Description

When you create SQL Macros, Vertica allows multiple user-defined functions to share the same name with different argument types. Therefore, if you try to alter or drop a function without specifying the argument data type, the system returns the following error message:

```
=> DROP FUNCTION my_function();  
ROLLBACK: Function with specified name and parameters does not exist: my_function
```

Resolution

Specify the argument type when you alter or drop a SQL Macro function; for example:

```
=> DROP FUNCTION my_function(x INT);
```

Tip: To view a list of all SQL Macro functions on which you have EXECUTE privileges, (which returns their argument types), query the V_CATALOG.USER_FUNCTIONS system table.

See the following topics in the SQL Reference Manual:

- CREATE FUNCTION
- ALTER FUNCTION
- DROP FUNCTION
- USER_FUNCTIONS

Hash Table Out of Memory on COPY and REFRESH Operations

Error

```
ERROR: Join inner did not fit in memory
ERROR: HASH JOIN inner did not fit in memory
ERROR: Hash table out of memory
```

Description

A hash join table that is too large to fit in memory triggers an error that can occur on joins that execute using a hash join (rather than a merge join) where the inner relation is large, such as a dimension table or outer join and while loading data into a pre-join projection for COPY and REFRESH operations

Resolution

1 Enable join spill by trying one or more of the following:

- To enable join spilling for a session:

```
=> SELECT add_vertica_options('EE', 'ENABLE_JOIN_SPILL');
```

- To disable join spilling for a session:

```
=> SELECT clr_vertica_options('EE', 'ENABLE_JOIN_SPILL');
```

- To enable join spilling for a single query, use hint syntax:

```
=> SELECT /*+set_vertica_options(EE, ENABLE_JOIN_SPILL)*/ <rest of the query>;
```

Only one hint can be used in a statement. For example, in an "INSERT /*+direct*/ SELECT ..." query where the select list includes a join that requires ENABLE_JOIN_SPILL, set ENABLE_JOIN_SPILL with the following:

```
=> SELECT add_vertica_options('EE', 'ENABLE_JOIN_SPILL'); and
```

```
=> SELECT clr_vertica_options('EE', 'ENABLE_JOIN_SPILL')
```

- To globally enable this option for all new sessions started in the database:

```
=> SELECT
set_config_parameter('EEVerticaOptions', 'ENABLE_JOIN_SPILL');
```

IMPORTANT: Setting ENABLE_JOIN_SPILL globally or leaving it on can degrade performance for subsequent queries, especially for queries that do not need it. Thus, Vertica recommends that you enable the join spill mechanism only when needed and then disable it immediately after the query successfully runs.

2 Re-execute the query.

I/O Error

Description

Under some circumstances, an I/O error caused by a failed disk could cause the server to terminate with a SIGBUS error.

Resolution

If a bus error is received, check the operating system logs, such as *dmesg*, to see if information about a hardware error is present.

Insufficient Projections to Answer Query

Error

```
Insufficient Projections to answer query
```

Description

Possible causes:

- 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.

Resolution

- 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

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

Description

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

Resolution

None

Loading Data While a Node is Down

You might see the following warning while executing a COPY command:

Error

WARNING: the number of rows reported as loaded may be incorrect if a node is down.

Description

If you load data while a node is down, the number of rows reported as loaded in vsql is incorrect. The *vertica.log* file correctly reports the actual number of rows loaded, which can also be confirmed through `SELECT COUNT (*)` on the table being loaded.

Load Stream Could Appear Complete While Sorting in Progress

Description

When loading data, the load stream appears complete even though the data is still being sorted. VT_STREAMS table might show 100% completed as it is tracking the data parsed and does not account for sorting.

Lock Timeout Error Reports Incorrect Transaction Isolation Level

Error

```
ERROR: Locking failure: Timed out I locking Table:tab. S held by  
[user release (select min(col) from tab;)]. Your current transaction isolation  
level is SERIALIZABLE
```

Description

If you set the session characteristics as `TRANSACTION ISOLATION LEVEL READ COMMITTED` and execute some DDL or run refresh or Tuple Mover operations, you might see the above error, which shows your isolation setting to be different from what you set.

Vertica internal processes and DDL operations are run at `SERIALIZABLE` isolation to ensure consistency.

Resolution

None. Behaves as expected. To confirm your sessions transaction isolation level, use the `SHOW` command:

```
=> SHOW TRANSACTION ISOLATION LEVEL;
```

Missing FROM-clause Entry for Table

Description

An alias is specified for a table but not used in the WHERE clause gives an error.

Example:

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

Resolution

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;
```

Operator Does Not Exist

Issue

Vertica does not automatically convert date type values to a character string when they are used in an expression. For example, the following query, which tries to return the `date_column` column surrounded by spaces and single quotes, returns an error.

```
SELECT ' ' || date_column || ' ' from T1;
ERROR: operator does not exist: "unknown" || date
HINT: No operator matches the given name and argument type(s). You may need to add explicit type casts.
```

Resolution

Explicitly cast the `date_column` column to a character string:

```
SELECT ' ' || date_column::varchar || ' ' from T1;
```

Out of Memory

Error

```
Out of memory
```

Description

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

Resolution

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

Partitioning Errors

The following error messages could occur when you trying to partition tables.

Error

```
ROLLBACK: Table is used as a dimension in a prejoined projection
```

Description

DROP_PARTITION fails if the table specified is used as a dimension table in a pre-join projection

Error

```
Too many ROS containers
```

Description

If the table has too many partitions, Vertica can return a "Too many ROS containers" message.

Resolution

- 1 Create a new table with proper partitioning.
- 2 Reload the data using INSERT...SELECT operations.
- 3 Drop the original table.

You could also run DO_TM_TASK to perform a Tuple Mover operation (moveout) on one or more projections defined on a specified table.

Alternatively, you can perform a manual mergeout using ALTER PROJECTION mergeout, but if the partitioned data is too granular you could encounter the issue again. Therefore, it is best to plan by defining fewer than 12 partitions. Otherwise, restructure the table or contact **Technical Support** (on page 1) if you encounter errors.

Caution: ALTER PROJECTION mergeout and PURGE() operations ignore partitions and partition boundaries while merging ROS containers. So after a manual mergeout/purge, DROP_PARTITION could cause the system to force partition the data. Vertica recommends that you use DO_TM_TASK(moveout) instead.

See Also

Cannot drop partition with nodes down (page 71)

Managing Large Grouped ROS Containers (page 46)

CREATE PROJECTION fails if projection excludes columns needed for partitioning (page 81)

ALTER PROJECTION, DROP_PARTITION, PARTITION_TABLE, and PARTITIONS in the SQL Reference Manual

Partitioning Tables in the Administrator's Guide

Projection Errors

Case 1

The following error message can occur when you try to create projections.

- **CREATE PROJECTION** fails if the projection does not include the columns it needs for partitioning and returns a message:
ROLLBACK: Attribute "month" in the projection anchor table "fact" must be included in the projection
Vertica recommends that you load data on a partition-by-partition basis to avoid the additional step of force partitioning the data later on.

Case 2

The following error messages can occur when you try to drop projections.

- If the projection you want to drop is the projection with the longest history, Vertica returns a message:
drop projection mytable_p1;
ROLLBACK: Projection cannot be dropped because history after AHM would be lost
DETAIL: Current AHM epoch = 10
To resolve this issue, use **MAKE_AHM_NOW**, which sets the Ancient History Mark to the greatest allowable value and lets you drop any projections that existed before the issue occurred.
- If the projection you want to drop is the last superprojection on the table, Vertica returns a message:
drop projection A_p2;
ROLLBACK: No up-to-date super projection left on the anchor table of projection a_p2
HINT: Use **DROP TABLE ... CASCADE** to drop the anchor table and its last projection, or create a replacement super projection instead
To resolve this issue, drop the anchor table using the **CASCADE** keyword or create another superprojection.

See Also

DROP_PARTITION in the SQL Reference Manual

Partitioning Tables in the Administrator's Guide

Projection Refresh Doesn't Start Until a Buddy Projection is Created

Error

Starting refresh background process

Description

A projection is not refreshed until after a buddy projection is created.

After CREATE PROJECTION is executed, if you execute `'select start_refresh()'`, Vertica returns an error.

The refresh, however, does not actually begin until after a buddy projection is created.

Resolution

You can monitor the refresh operation by examining the `vertica.log` file or view the final status of the projection refresh by using:

```
=> SELECT GET_PROJECTIONS('table_name')
```

Vertica internal operations (mergeout, refresh, and recovery) maintain partition separation except in certain cases:

- Recovery of a projection when the buddy projection from which the partition is recovering is identically sorted. If the projection is undergoing a full rebuild, it is recovered one ROS container at a time. The projection ends up with a storage layout identical to its buddy and is, therefore, properly segmented.

Note: In the case of a partial rebuild, all recovered data goes into a single ROS container and must be partitioned manually.

- Manual tuple mover operations often output a single storage container, combining any existing partitions; for example, after executing any of the `PURGE()` operations.

Optimizer Does Not Select the Best Projection

Error

EXPLAIN plan could show that the optimizer is not selecting the best projection for a query.

Resolution

Run `ANALYZE_STATISTICS()` for the tables involved:

```
=> SELECT ANALYZE_STATISTICS('table_name');
```


Query Continues to Execute on the Server When the Client is Killed

Description

When a vsql or JDBC client is killed or crashes, any executing query continues to execute on the server until it attempts to communicate with the client.

Resolution

You can terminate the server portion of the query using the INTERRUPT_STATEMENT function in the Session Management API.

See Also

INTERRUPT_STATEMENT in the SQL Reference Manual

Query Required More Resources Than Initiator Resource Manager Estimated

Error

```
Query required more resources than initiator resource manager estimated,  
likely due to high storage container counts or a heavier workload on node <...>
```

Description

Heavy data load conditions can cause the Tuple Mover to fall behind in performing moveout or mergeout operations. The resulting large number of ROS containers can cause some requests to exhaust all available system resources. Vertica detects this problem and prevents load transactions until the Tuple Mover has time to catch up.

Resolution

You might need to adjust the Tuple Mover's configuration parameters to compensate for the load pattern or rate. See *Tuning the Tuple Mover* in the Administrator's Guide for details.

You can query the TUPLE_MOVER_OPERATIONS table to monitor mergeout activity. However, the Tuple Mover does not immediately start a mergeout when a projection reaches the limit of ROS containers, so you may not see a mergeout in progress when receiving this error.

If waiting for a mergeout does not resolve the error, the problem probably is related to insufficient RAM. In that case, contact **Technical Support** (on page 1). A good rule of thumb is that system RAM in MB divided by 6 times the number of columns in the largest table should be greater than 10. For example, for a 100 column table you would want at least 6GB of RAM ($6144\text{MB} / (6 * 100) = 10.24$) to handle continuous loads.

Querying Monitoring Table When a Node is Recovering

Description

Currently querying the monitoring tables on a node that is recovering is not supported.

Record Terminator for COPY When Loading Data from a Windows Client

Description

The default record terminator for COPY is '\n'. If you are loading data from a Windows client, you must specify RECORD TERMINATOR '\r\n'.

If you are using JDBC, Vertica recommends that you use the value from

```
System.getProperty("line.separator")
```

for the RECORD TERMINATOR

Referential Integrity Load Violation

Description

Vertica checks for constraint violations when queries are executed, not when data is loaded.

If you have a pre-joined projection defined on the table being loaded, Vertica checks for constraint violations (duplicate primary keys or non-existent foreign keys) during the join operation and reports errors. If there are no pre-joined projections, Vertica performs no such checks.

To avoid constraint violations, load data without committing it and then perform a post-load check of your data using the ANALYZE_CONSTRAINTS function. If the function finds constraint violations, you can roll back the load because you have not committed it.

See Also

Analyzing Constraints (Detecting Constraint Violations) in the Administrator's Guide

COPY NO COMMIT and ANALYZE_CONSTRAINTS in the SQL Reference Manual

Too Many ROS Containers

Error

```
Too many ROS containers exist for the following projections:  
toomanyros_projection(969 ROSs).  
Please wait for the tuple mover to catch up. Use 'select *  
from tuple_mover_operations' to monitor.
```

Description

Heavy data load conditions can cause the Tuple Mover to fall behind in performing moveout or mergeout operations. The resulting large number of ROS containers can cause some requests to exhaust all available system resources. Vertica detects this problem and prevents load transactions until the Tuple Mover has time to catch up.

Resolution

You might need to adjust the Tuple Mover's configuration parameters to compensate for the load pattern or rate. See [Tuning the Tuple Mover in the Administrator's Guide](#) for details.

You can query the `TUPLE_MOVER_OPERATIONS` table to monitor mergeout activity. However, the Tuple Mover does not immediately start a mergeout when a projection reaches the limit of ROS containers, so you may not see a mergeout in progress when receiving this error.

If waiting for a mergeout does not resolve the error, the problem probably is related to insufficient RAM. In that case, contact [Technical Support](#) (on page 1). A good rule of thumb is that system RAM in MB divided by 6 times the number of columns in the largest table should be greater than 10. For example, for a 100 column table you would want at least 6GB of RAM ($6144\text{MB} / (6 * 100) = 10.24$) to handle continuous loads.

Update and Delete Performance

Description

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

Resolution

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

See Also

[Best practices for DELETE and UPDATE in Administrator's Guide](#)

Using Stream Name Parameter while Loading Data from STDIN

Description

When using COPY to load data from STDIN, always use the `STREAM NAME` parameter. This lets you distinguish between multiple loads from STDIN when you view the `VT_LOAD_STREAMS` table, and also lets you easily associate messages in the log file with specific data loads.

Variable Hash Join Performance

Description

A hash join to an inner relation could require additional memory than was initially allocated to the query. If memory is available, the hash join executes the query without needing to spill to disk. If memory is not available, the query automatically re-executes with join spill enabled (ENABLE_JOIN_SPILL), but query performance will be slower.

Resolution

If you notice performance degradation or variable performance for a hash join, try one of the following:

- Confirm that all tables involved in the join or joins have statistics:
=> `SELECT * FROM PROJECTIONS WHERE NOT has_statistics;`
- Look in the `vertica.log` for the following messages:

When the query is planned and has no statistics for projections used:

```
2010-08-15 04:19:57.824 Init Session:0xc2972d0 [Command]
<INFO> Query being replanned with ENABLE_JOIN_SPILL used
projections with no statistics: lineorder_node01 cust_p1
2010-08-15 04:19:57.824 Init Session:0xc2972d0 [Command]
<INFO> Query Retry: Setting
    add_vertica_options('EE', 'ENABLE_JOIN_SPILL');
```

Note: If the table has statistics, the optimizer might chose a different inner table for the join.

When the query is planned and has statistics for projections used:

```
2010-08-15 04:19:57.824 Init Session:0xc2972d0 [Command]
<INFO> Query Retry: Setting
    add_vertica_options('EE', 'ENABLE_JOIN_SPILL');
```

The following message displays when a query cannot be replanned; for example, when query results are output on one node before another node runs out of memory for its hash table. This can happen when there is a data skew (one node has more data than another or a subquery has differing amounts of data on different nodes):

```
2010-08-16 16:37:39.170 Init Session:0x2aaab4019190 [Command]
<INFO> Query can't be replanned due to partial output from initial
    execution
```

If you want more predictable results:

- 1 Profile the query to look for the minimum memory required for the query to run:

```
=> PROFILE SELECT ... FROM ...;
```

A profiled query returns the following message before the results:

```
NOTICE: Initiator memory estimate for query: [on pool general: 22660
    KB,
    minimum: 22660 KB]
```

- 2 Create a standalone resource pool, which specifies at a minimum the amount of memory that the PROFILE command showed; for example, the following resource pool allocates 30MB of memory:

```
CREATE RESOURCE POOL <pool_name> MEMORYSIZE '30M';
```

- 3 Check `vertica.log` and to make sure there is no message about a query retry, such as in the following fragment:

```
2010-08-16 16:37:39.170 Init Session:0x2aaab4019190 [Command] <INFO>  
Query can't be replanned due to partial output from initial execution
```

See Also

CREATE RESOURCE POOL and Built-in Pools in the SQL Reference Manual

Platform Problems

This section describes platform-specific issues you might encounter.

Data Disk Spontaneously Remounts

Description

A temporary data storage disk using an ext3 file system could 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

To increase free memory perform the following steps:

- 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`

Setting the Device Readahead Parameter for Improved Performance

Error

During installation you might see

```
Readahead parameter is set too high for some devices on the host
  Consider adding the following lines to /etc/rc.local file for better
performance
  blockdev --setra 2048 /dev/sda2 # old value 12024
```

Description

The default device readahead parameter could be too high, resulting in contention for disk cache space and slowing down the `mergeout` process. If you see the above message during the installation, Vertica strongly recommends that you follow the instructions to add the `blockdev` command to the `/etc/rc.local` file.

Move Core File to a Specified Location

Explanation

To control where core files are placed, explicitly specify their location by setting `/proc/sys/kernel/core_pattern`.

Note: This is a SYSTEM WIDE setting that requires root access to change. The precise format available can be retrieved from "man proc" or at this web page:
<http://www.ncsa.uiuc.edu/UserInfo/Resources/Hardware/CommonDoc/core.html>

By default, the `core_pattern` is just "core," which results in a file named `core.pid`. This file resides in the current working directory, which by default, is the catalog directory for Vertica.

SUSE 10/11 Install Fails During Spread Configuration

Description

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

Resolution

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.

SSL Problems

This section describes SSL issues you might encounter.

Unsafe permissions on private key file

Error

Unsafe permissions on private key file <filename>. File must be owned by the database user and must have no permissions for "group" or "other"

Description

The permissions for the specified key file are too permissive. This is a fatal error, and the database fails to start.

Resolution

Remove any "group" or "other" permissions for the specified file.

Could not access server public certificate file

Error

Could not access server public certificate file "<filename>": No such file or directory

Description

The SSL stack could not be initialized because server.crt is missing or corrupted. This is a fatal error, and the database fails to start.

Resolution

See SSL Prerequisites and verify that a valid server certificate file is located in the appropriate directory.

Could not access server private key file

Error

Could not access server private key file: <detail>

Description

The SSL stack could not be initialized because server.key is missing or corrupted. This is a fatal error, and the database fails to start.

Resolution

See SSL Prerequisites and verify that a valid server key file is located in the appropriate directory.

Check of private key failed

Error

Check of private key failed: <error_detail>

Description

The public certificate and private key file are not a pair. If, for example, the public certificate is from key pair A and the private key is from key pair B, they do not form a pair. This is a fatal error, and the database fails to start.

Resolution

See SSL Prerequisites to set up an appropriate public certificate and private key file pair.

Virtual IP Server Problems

This section describes issues you might encounter related to Vertica Load Balancer. See Also Load Balancing in the Administrator's Guide.

Users Cannot Connect to the VIP

Issue

Users cannot connect to the database through the Virtual IP (VIP) address.

Resolution

1 Check if spread is running:

```
$ ps ax | grep spread
11895 ?          S<s      4:30 /opt/vertica/spread/sbin/spread -n
      N192168051055 -c
                                /opt/vertica/config/vspread.conf
29617 pts/3      S+       0:00 grep spread
```

1. If spread is not running, start spread as root or using sudo:

```
[root@node01]# /etc/init.d/spread start
```

2. If spread is running, restart spread as root or using sudo:

```
[root@node01]# /etc/init.d/spread restart
```

3. Check the spread status as root or using sudo:

```
[root@node01]# /etc/init.d/spread status
```

4. Issue the `ifconfig` command to check the current IP addresses of the hosts, and verify that those IP addresses are listed in `/opt/vertica/config/vspread.conf`.

```
[root@node01]# ifconfig
```

If spread fails to start, examine the following files for problems:

```
/tmp/spread*.log
```

```
/var/log/spread.log
```

Permission problems and syntax problems are identified in the log files.

2 Check if keepalived is running:

```
$ ps ax | grep keepalived
29622 pts/3      S+       0:00 grep keepalived
```

1. If keepalived is not running, start keepalived as root or using sudo:

```
# /etc/init.d/keepalived start
```

2. If keepalived is running, restart keepalived as root or using sudo:

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

Issue

Users cannot connect to the database.

Resolution

Try to telnet to the VIP and port:

```
# telnet 10.10.51.180 5433
```

If telnet reports no route to host, recheck your `/etc/keepalived/keepalived.conf` file to make sure you entered the correct VIP and RIPs.

Errors and informational messages from the keepalived daemon are written to the `/var/log/messages` file, so check the messages file first:

```
# tail -f /var/log/messages
May 18 09:04:32 dell02 Keepalived_vrrp: VRRP_Instance(VI_1) Sending gratuitous
ARPs on eth0 for 10.10.10.100
May 18 09:04:32 dell02 avahi-daemon[3191]: Registering new address record for
10.10.10.100 on eth0.
May 18 09:04:32 dell02 Keepalived_healthcheckers: Netlink reflector reports IP
10.10.10.100 added
```

Expected e-mail messages from the keepalived daemon

- Upon startup:
Subject: [node01] VRRP Instance VI_1 - Entering MASTER state
=> VRRP Instance is now owning VRRP VIPs <=
- When a node fails:
Subject: [node01] Realserver 10.10.10.1:5433 - DOWN
=> MISC CHECK failed on service <=
- When a node comes back up:
Subject: [node02] Realserver 10.10.10.1:5433 - UP

=> MISC CHECK succeed on service <=

Resolving Keepalived Issues

If there are connection or other issues related to the Virtual IP server and Keepalived, try some of the following tips:

- Set `KEEPALIVED_OPTIONS="-D -d"` in the `/etc/sysconfig/keepalived` file to enable both debug mode and dump configuration.
- Monitor the system log in `/var/log/messages`. If `keepalived.conf` is incorrect, the only indication is in the messages log file. For example:
\$ tail /var/log/messages
Errors and informational messages from the keepalived daemon are also written to the `/var/log/messages` files.
- Type `ip addr list` and see the configured VIP addresses for eth0. For example:
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 16436 qdisc noqueue
link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00

```
inet 127.0.0.1/8 scope host lo
inet 10.10.51.180/32 brd 127.255.255.255 scope global lo:0
inet6 ::1/128 scope host
    valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast qlen
1000
    link/ether 84:2b:2b:55:4b:be brd ff:ff:ff:ff:ff:ff
    inet 10.10.51.55/24 brd 10.10.51.255 scope global eth0
    inet6 fe80::862b:2bff:fe55:4bbe/64 scope link
    valid_lft forever preferred_lft forever
3: eth1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast
qlen 1000
    link/ether 84:2b:2b:55:4b:bf brd ff:ff:ff:ff:ff:ff
    inet 192.168.51.55/24 brd 192.168.51.255 scope global eth1
    inet6 fe80::862b:2bff:fe55:4bbf/64 scope link
    valid_lft forever preferred_lft forever
4: sit0: <NOARP> mtu 1480 qdisc noop
    link/sit 0.0.0.0 brd 0.0.0.0
```

- Check `iptables` and notice the `PREROUTING` rule on the `BACKUP` (slave) director. Even though `ipvsadm` has a complete list of real servers to manage, it does not route anything as the prerouting rule redirects packets to the loopback interface.

```
# /sbin/iptables -t nat -n -L
Chain PREROUTING (policy ACCEPT)
target     prot opt source                destination
Chain POSTROUTING (policy ACCEPT)
target     prot opt source                destination
Chain OUTPUT (policy ACCEPT)
target     prot opt source                destination
```

Note: On some kernels, the `nat` tables does not show by default without the `-t` parameter, and `-n` is used to avoid long DNS lookups. See the *iptables(8) - Linux man page* <http://linux.die.net/man/8/iptables> for details.

- During failover, it is normal to expect delay in new connection establishment until the slave node takes control. The delay could be several minutes depending on the load on the cluster. If you cannot connect to the database, try to telnet to the VIP and port:

```
# telnet 10.10.51.180 5433
```

If telnet reports no route to host, recheck the `keepalived` configuration file (`/etc/keepalived/keepalived.conf`) to make sure you entered the correct VIP and RIPv.

Index

(

(LOSTCONTACT) Slow Database Startup when Node is Down • 65

.

... large moveout is in progress • 31

... users are connected • 32, 62

A

About the Documentation • 2

Administration Tools Errors • 52

Administration Tools shows node state as UNKNOWN • 40

Attempted to Create Too Many ROS Containers • 72

B

Backup and Restore Problems • 50

C

Cannot Allocate Memory • 52

Cannot be Retired • 70

Cannot bootstrap catalog on top of existing catalog • 50

Cannot Drop Partition with Nodes Down • 72, 81

Cannot find IP Address for Host... • 52

Case-sensitive System Table Queries Return 0 Records • 73

Catalog

 Dumping • 26

Check of private key failed • 92

Clearing Profiling Data • 16

Clock Skew Detected • 53

Collecting Query Information • 17, 18

Column Definition Changed to NOT NULL • 73

Configuration Problems • 52

Configuring Query Repository • 23, 24

Connection Problems • 61

COPY, DELETE, UPDATE, or INSERT Hangs When Server Idle • 74

Copyright Notice • 100

Could not access server private key file • 91

Could not access server public certificate file • 91

Could Not Add Location (Permission) • 67

Could Not Add Location (Unavailable) • 68

Could Not Connect to Database • 61

Could Not Connect to Server

 Connection Refused • 61

Could Not Create Directory • 67

Could Not Log in with SSH • 53

Could Not Open Configuration File • 54

Could Not Open Directory • 67

Could Not Write To Filename • 74

Create Database Fails • 54

Ctrl-C Distorts Query Result Set Output • 74

D

Data Disk Spontaneously Remounts • 89

Database Designer Errors • 64

Database did not start cleanly on initiator node! • 37

Database startup successful, but it could be incomplete • 36

dbLog • 38

Deleting Records from Table with PK with Associated FK • 75

Denied. Not a super user. Cannot create a schema • 54

Determining the Version Number • 27

Diagnosing spread problems • 40

Diagnostics • 1, 26, 27

Diagnostics Utility • 1, 27, 44

Disabling Profiling for All Sessions • 12

Disabling Profiling for the Current Session • 12

DROP_LOCATION for Data Locations is not Supported • 70

DUMP_CATALOG • 26

E

ENABLE_JOIN_SPILL • 77

Enabling and Disabling Profiling • 11, 13

Enabling and Disabling Query Repository • 18, 20

Enabling Profiling for All Sessions • 12

Enabling Profiling for the Current Session • 11

EOF error • 45, 46

Error Detected During Node Verification • 55

Error starting database, no nodes are up • 36

Errors found in SAL or Catalog • 50

Errors that Result in a Statement-Level Rollback • 75

Errors that Result in a Transaction-Level Rollback • 75

etc/hosts • 52

Execution Error on External Procedures • 75

Execution Problems • 72
Exporting a Catalog • 28

F

Failed Dependencies • 43
Failure Recovery • 29
Firewall Configuration • 62
Firewall inactivity timeout hit (Insufficient
KEEPALIVE value) • 63
Function <function-name>(<data-type>, ...) does
not exist • 76
Function with specified name and parameters
does not exist • 76

G

Good epoch logs are available on all nodes • 37

H

Hash Table Out of Memory on COPY and
REFRESH Operations • 77
Host Key Verification Failure • 46

I

I/O Error • 78
 Permission Denied • 58
Installation Problems • 43
Installation runs as sudoer when run using sudo
root • 47
Insufficient Privilege USAGE on SCHEMA • 55
Insufficient Projections to Answer Query • 78

K

KeyError
 'getpwuid()
uid not found
 <uid>' • 65

L

Latency Should be Greater Than (>) One (1) • 70
LIKE Predicate Referred to as ~• • 78
Load Stream Could Appear Complete While
 Sorting in Progress • 79
Loading Data While a Node is Down • 78
Location Cannot be Dropped • 71
Lock table • 26
Lock Timeout Error Reports Incorrect
 Transaction Isolation Level • 79

M

Managing And Viewing Query Repository • 24
Managing Large Grouped ROS Containers • 47,
 81
Misleading Message Due to IP/hostname
 Mismatch • 45
Missing FROM-clause Entry for Table • 79
Monitoring the QUERY_REPO Table • 18
Move Core File to a Specified Location • 90

N

New Session Rejected Due to Limit • 62
No good epoch log available on node • 38
No running statement, that session is idle • 35
No Such File or Directory ...
 super_projections.tmp.sql • 64
Node does not recover because of lock timeouts •
 39
Nodes stuck in INITIALIZING state • 38
Not Enough Open File Handles Allowed... • 55

O

One or More Nodes Did Not Open a Data
 Connection to this Node • 59
Operator Does Not Exist • 80
Optimizer Does Not Select the Best Projection •
 83
Out of Memory • 80

P

pam_limits • 65
Partitioning Errors • 81
Passwordless SSH Access to Other Sites Failed •
 45
Permission denied even after providing correct
 root password • 46
Permission Denied for Relation • 59
Platform Problems • 89
Preface • 9
Printing Full Books • 4
Profiling a Single Statement • 13
Profiling Database Performance • 11, 18, 19, 23
Projection Errors • 81, 82
Projection Refresh Doesn't Start Until a Buddy
 Projection is Created • 83

Q

- Query Continues to Execute on the Server When the Client is Killed • 84
- Query repository • 18, 23, 24
- Query Required More Resources Than Initiator Resource Manager Estimated • 84
- Querying Monitoring Table When a Node is Recovering • 85

R

- Reading the Online Documentation • 2
- Real-time Profiling • 13
- Record Terminator for COPY When Loading Data from a Windows Client • 85
- Referential Integrity Load Violation • 85
- Resolving Keepalived Issues • 94
- Restore fails if files present in catalog directory • 51
- Running Diagnostics Utility for Failed Installation • 27, 44

S

- Setting the Device Readahead Parameter for Improved Performance • 89
- Shutdown Problems • 31
- Spread Dead but pid File Exists • 59
- Spread is not running • 39, 40
- Spread Problems • 39
- SSH Already Configured • 44
- SSL • 91, 92
- SSL Problems • 91
- Startup Problems • 31, 35, 65
- Startup successful, but some nodes are recovering • 35
- Storage Location Problems • 67
- Suggested Reading Paths • 2, 4
- SUSE 10/11 Install Fails During Spread Configuration • 90

T

- Technical Support • 1, 4, 36, 37, 38, 39, 40, 41, 44, 47, 49, 72, 81, 84, 86
- Throughput Should be Greater Than (>) Zero (0) • 69
- TIMEOUT ERROR
 - Could not login with SSH • 37
- Too Many ROS Containers • 86

- Troubleshooting Overview • 10
- Typographical Conventions • 7

U

- Unable to drop table. Must be Owner of Relation • 58
- Unsafe permissions on private key file • 91
- Update and Delete Performance • 86
- Usage Cannot be an Empty String • 68
- Usage Cannot be Changed to Data • 68
- Usage Cannot be Changed to Temp • 69
- Usage of Storage Location Cannot be Changed • 69
- Users Cannot Connect to the VIP • 93
- Using Diagnostic Tools • 26, 38
- Using Stream Name Parameter while Loading Data from STDIN • 87
- Using sudo NOFILE Parameters Not Saved in limits.conf • 57

V

- Variable Hash Join Performance • 87
- Vertica Fails to Start (verticad and pam_limits) • 65
- verticad • 65
- Viewing Profiling Data • 15
- Viewing Real-time Profiling Data • 16
- Virtual IP Server Problems • 93

W

- Where to Find Additional Information • 6
- Where to Find the Vertica Documentation • 2

Copyright Notice

Copyright© 2006-2011 Vertica Systems, Inc., and its licensors. All rights reserved.

Vertica Systems, Inc.
8 Federal Street
Billerica, MA 01821
Phone: (978) 600-1000
Fax: (978) 600-1001
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. 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™, the Vertica® Analytic Database™, and FlexStore™ 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.

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.

Open Source Software Acknowledgments

Vertica makes no representations or warranties regarding any third party software. All third-party software is provided or recommended by Vertica on an AS IS basis.

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com).

ASMJIT

Copyright (c) 2008-2010, Petr Kobalicek <kobalicek.petr@gmail.com>

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

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 AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Boost

Boost Software License - Version 1.38 - February 8th, 2009

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.

bzip2

This file is a part of bzip2 and/or libbzip2, a program and library for lossless, block-sorting data compression.

Copyright © 1996-2005 Julian R Seward. All rights reserved.

- 1** Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:
- 2** Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- 3** The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.
- 4** Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
- 5** The name of the author may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Julian Seward, Cambridge, UK.

jseward@bzip.org <<mailto:jseward@bzip.org>>

bzip2/libbzip2 version 1.0 of 21 March 2000

This program is based on (at least) the work of:

Mike Burrows

David Wheeler

Peter Fenwick

Alistair Moffat

Radioed Neal

Ian H. Witten

Robert Sedgewick

Jon L. Bentley

Daemonize

Copyright © 2003-2007 Brian M. Clapper.

All rights reserved.

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

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the clapper.org nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Ganglia Open Source License

Copyright © 2001 by Matt Massie and The Regents of the University of California.

All rights reserved.

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

IN NO EVENT SHALL THE UNIVERSITY OF CALIFORNIA BE LIABLE TO ANY PARTY FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES 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 OBLIGATION TO PROVIDE MAINTENANCE, SUPPORT, UPDATES, ENHANCEMENTS, OR MODIFICATIONS.

ICU (International Components for Unicode) License - ICU 1.8.1 and later

COPYRIGHT AND PERMISSION NOTICE

Copyright © 1995-2009 International Business Machines Corporation and others

All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, provided that the above copyright notice(s) and this permission notice appear in all copies of the Software and that both the above copyright notice(s) and this permission notice appear in supporting documentation.

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 AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR HOLDERS INCLUDED IN THIS NOTICE BE LIABLE FOR ANY CLAIM, OR ANY SPECIAL INDIRECT OR CONSEQUENTIAL DAMAGES, OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Except as contained in this notice, the name of a copyright holder shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization of the copyright holder.

All trademarks and registered trademarks mentioned herein are the property of their respective owners.

Keepalived Vertica IPVS (IP Virtual Server) Load Balancer

Copyright © 2007 Free Software Foundation, Inc.

<http://fsf.org/>

The keepalived software contained in the `VerticaIPVSLoadBalancer-4.1.x.RHEL5.x86_64.rpm` software package is licensed under the GNU General Public License ("GPL"). You are entitled to receive the source code for such software. For no less than three years from the date you obtained this software package, you may download a copy of the source code for the software in this package licensed under the GPL at no charge by visiting <http://www.vertica.com/licenses/keepalived-1.1.17.tar.gz> <http://www.vertica.com/licenses/keepalived-1.1.17.tar.gz>. You may download this source code so that it remains separate from other software on your computer system.

jQuery

Copyright © 2009 John Resig, <http://jquery.com/>

All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

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 AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Lighttpd Open Source License

Copyright © 2004, Jan Kneschke, incremental

All rights reserved.

- 1** Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:
- 2** Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- 3** Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- 4** Neither the name of the 'incremental' nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

MersenneTwister.h

Copyright © 1997 - 2002, Makoto Matsumoto and Takuji Nishimura,

Copyright © 2000 - 2009, Richard J. Wagner

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.
- 2** Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- 3** The names of its contributors may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

MIT Kerberos

Copyright © 1985-2007 by the Massachusetts Institute of Technology.

Export of software employing encryption from the United States of America may require a specific license from the United States Government. It is the responsibility of any person or organization contemplating export to obtain such a license before exporting.

WITHIN THAT CONSTRAINT, permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of M.I.T. not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. Furthermore if you modify this software you must label your software as modified software and not distribute it in such a fashion that it might be confused with the original MIT software. M.I.T. makes no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

Individual source code files are copyright MIT, Cygnus Support, Novell, OpenVision Technologies, Oracle, Red Hat, Sun Microsystems, FundsXpress, and others.

Project Athena, Athena, Athena MUSE, Discuss, Hesiod, Kerberos, Moira, and Zephyr are trademarks of the Massachusetts Institute of Technology (MIT). No commercial use of these trademarks may be made without prior written permission of MIT.

"Commercial use" means use of a name in a product or other for-profit manner. It does NOT prevent a commercial firm from referring to the MIT trademarks in order to convey information (although in doing so, recognition of their trademark status should be given).

Portions of src/lib/crypto have the following copyright:

Copyright © 1998 by the FundsXpress, INC.

All rights reserved.

Export of this software from the United States of America may require a specific license from the United States Government. It is the responsibility of any person or organization contemplating export to obtain such a license before exporting.

WITHIN THAT CONSTRAINT, permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of FundsXpress. not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. FundsXpress makes no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

THIS SOFTWARE IS PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

The implementation of the AES encryption algorithm in src/lib/crypto/aes has the following copyright:

Copyright © 2001, Dr Brian Gladman <brg@gladman.uk.net>, Worcester, UK.
All rights reserved.

LICENSE TERMS

The free distribution and use of this software in both source and binary form is allowed (with or without changes) provided that:

- 1 Distributions of this source code include the above copyright notice, this list of conditions and the following disclaimer.
- 2 Distributions in binary form include the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other associated materials.
- 3 The copyright holder's name is not used to endorse products built using this software without specific written permission.

DISCLAIMER

This software is provided 'as is' with no explicit or implied warranties in respect of any properties, including, but not limited to, correctness and fitness for purpose.

The implementations of GSSAPI mechglue in GSSAPI-SPNEGO in src/lib/gssapi, including the following files:

- lib/gssapi/generic/gssapi_err_generic.et
- lib/gssapi/mechglue/g_accept_sec_context.c
- lib/gssapi/mechglue/g_acquire_cred.c
- lib/gssapi/mechglue/g_canon_name.c
- lib/gssapi/mechglue/g_compare_name.c
- lib/gssapi/mechglue/g_context_time.c
- lib/gssapi/mechglue/g_delete_sec_context.c
- lib/gssapi/mechglue/g_dsp_name.c
- lib/gssapi/mechglue/g_dsp_status.c
- lib/gssapi/mechglue/g_dup_name.c
- lib/gssapi/mechglue/g_exp_sec_context.c
- lib/gssapi/mechglue/g_export_name.c
- lib/gssapi/mechglue/g_glue.c
- lib/gssapi/mechglue/g_imp_name.c

- lib/gssapi/mechglue/g_imp_sec_context.c
- lib/gssapi/mechglue/g_init_sec_context.c
- lib/gssapi/mechglue/g_initialize.c
- lib/gssapi/mechglue/g_inquire_context.c
- lib/gssapi/mechglue/g_inquire_cred.c
- lib/gssapi/mechglue/g_inquire_names.c
- lib/gssapi/mechglue/g_process_context.c
- lib/gssapi/mechglue/g_rel_buffer.c
- lib/gssapi/mechglue/g_rel_cred.c
- lib/gssapi/mechglue/g_rel_name.c
- lib/gssapi/mechglue/g_rel_oid_set.c
- lib/gssapi/mechglue/g_seal.c
- lib/gssapi/mechglue/g_sign.c
- lib/gssapi/mechglue/g_store_cred.c
- lib/gssapi/mechglue/g_unseal.c
- lib/gssapi/mechglue/g_userok.c
- lib/gssapi/mechglue/g_utils.c
- lib/gssapi/mechglue/g_verify.c
- lib/gssapi/mechglue/gssd_pname_to_uid.c
- lib/gssapi/mechglue/mglueP.h
- lib/gssapi/mechglue/oid_ops.c
- lib/gssapi/spnego/gssapiP_spnego.h
- lib/gssapi/spnego/spnego_mech.c

are subject to the following license:

Copyright © 2004 Sun Microsystems, Inc.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

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 AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Npgsql-.Net Data Provider for Postgresql

Copyright © 2002-2008, The Npgsql Development Team

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 NPGSQL DEVELOPMENT TEAM 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 NPGSQL DEVELOPMENT TEAM HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

THE NPGSQL DEVELOPMENT TEAM 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 NPGSQL DEVELOPMENT TEAM HAS NO OBLIGATIONS TO PROVIDE MAINTENANCE, SUPPORT, UPDATES, ENHANCEMENTS, OR MODIFICATIONS.

Open LDAP

The OpenLDAP Public License
Version 2.8, 17 August 2003

Redistribution and use of this software and associated documentation ("Software"), with or without modification, are permitted provided that the following conditions are met:

- 1** Redistributions in source form must retain copyright statements and notices,
- 2** Redistributions in binary form must reproduce applicable copyright statements and notices, this list of conditions, and the following disclaimer in the documentation and/or other materials provided with the distribution, and
- 3** Redistributions must contain a verbatim copy of this document.

The OpenLDAP Foundation may revise this license from time to time. Each revision is distinguished by a version number. You may use this Software under terms of this license revision or under the terms of any subsequent revision of the license.

THIS SOFTWARE IS PROVIDED BY THE OPENLDAP FOUNDATION AND ITS CONTRIBUTORS "AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OPENLDAP FOUNDATION, ITS CONTRIBUTORS, OR THE AUTHOR(S) OR OWNER(S) OF THE SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The names of the authors and copyright holders must not be used in advertising or otherwise to promote the sale, use or other dealing in this Software without specific, written prior permission. Title to copyright in this Software shall at all times remain with copyright holders.

OpenLDAP is a registered trademark of the OpenLDAP Foundation.

Copyright 1999-2003 The OpenLDAP Foundation, Redwood City, California, USA. All Rights Reserved. Permission to copy and distribute verbatim copies of this document is granted.

Open SSL

OpenSSL License

Copyright © 1998-2008 The OpenSSL Project. 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.
- 2 Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- 3 All advertising materials mentioning features or use of this software must display the following acknowledgment: "This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (<http://www.openssl.org/>)"
- 4 The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact openssl-core@openssl.org.
- 5 Products derived from this software may not be called "OpenSSL" nor may "OpenSSL" appear in their names without prior written permission of the OpenSSL Project.
- 6 Redistributions of any form whatsoever must retain the following acknowledgment: "This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>)"

THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT ``AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

PCRE LICENCE

PCRE is a library of functions to support regular expressions whose syntax and semantics are as close as possible to those of the Perl 5 language.

Release 8 of PCRE is distributed under the terms of the "BSD" licence, as specified below. The documentation for PCRE, supplied in the "doc" directory, is distributed under the same terms as the software itself.

The basic library functions are written in C and are freestanding. Also included in the distribution is a set of C++ wrapper functions.

THE BASIC LIBRARY FUNCTIONS

Written by: Philip Hazel
Email local part: ph10
Email domain: cam.ac.uk
University of Cambridge Computing Service,
Cambridge, England.
Copyright (c) 1997-2010 University of Cambridge
All rights reserved.

THE C++ WRAPPER FUNCTIONS

Contributed by: Google Inc.
Copyright (c) 2007-2010, Google Inc.
All rights reserved.

THE "BSD" LICENCE

- Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:
- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

Neither the name of the University of Cambridge nor the name of Google Inc. nor the names of their contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

End

Perl Artistic License

Copyright © August 15, 1997

Preamble

The intent of this document is to state the conditions under which a Package may be copied, such that the Copyright Holder maintains some semblance of artistic control over the development of the package, while giving the users of the package the right to use and distribute the Package in a more-or-less customary fashion, plus the right to make reasonable modifications.

Definitions

"Package" refers to the collection of files distributed by the Copyright Holder, and derivatives of that collection of files created through textual modification.

"Standard Version" refers to such a Package if it has not been modified, or has been modified in accordance with the wishes of the Copyright Holder as specified below.

"Copyright Holder" is whoever is named in the copyright or copyrights for the package.

"You" is you, if you're thinking about copying or distributing this Package.

"Reasonable copying fee" is whatever you can justify on the basis of media cost, duplication charges, time of people involved, and so on. (You will not be required to justify it to the Copyright Holder, but only to the computing community at large as a market that must bear the fee.)

"Freely Available" means that no fee is charged for the item itself, though there may be fees involved in handling the item. It also means that recipients of the item may redistribute it under the same conditions they received it.

- 1 You may make and give away verbatim copies of the source form of the Standard Version of this Package without restriction, provided that you duplicate all of the original copyright notices and associated disclaimers.
- 2 You may apply bug fixes, portability fixes and other modifications derived from the Public Domain or from the Copyright Holder. A Package modified in such a way shall still be considered the Standard Version.
- 3 You may otherwise modify your copy of this Package in any way, provided that you insert a prominent notice in each changed file stating how and when you changed that file, and provided that you do at least ONE of the following:
- 4 place your modifications in the Public Domain or otherwise make them Freely Available, such as by posting said modifications to Usenet or an equivalent medium, or placing the modifications on a major archive site such as uunet.uu.net, or by allowing the Copyright Holder to include your modifications in the Standard Version of the Package.
 1. use the modified Package only within your corporation or organization.
 2. rename any non-standard executables so the names do not conflict with standard executables, which must also be provided, and provide a separate manual page for each non-standard executable that clearly documents how it differs from the Standard Version.
 3. make other distribution arrangements with the Copyright Holder.
- 5 You may distribute the programs of this Package in object code or executable form, provided that you do at least ONE of the following:
 1. distribute a Standard Version of the executables and library files, together with instructions (in the manual page or equivalent) on where to get the Standard Version.
 2. accompany the distribution with the machine-readable source of the Package with your modifications.

3. give non-standard executables non-standard names, and clearly document the differences in manual pages (or equivalent), together with instructions on where to get the Standard Version.
4. make other distribution arrangements with the Copyright Holder.
- 6 You may charge a reasonable copying fee for any distribution of this Package. You may charge any fee you choose for support of this Package. You may not charge a fee for this Package itself. However, you may distribute this Package in aggregate with other (possibly commercial) programs as part of a larger (possibly commercial) software distribution provided that you do not advertise this Package as a product of your own. You may embed this Package's interpreter within an executable of yours (by linking); this shall be construed as a mere form of aggregation, provided that the complete Standard Version of the interpreter is so embedded.
- 7 The scripts and library files supplied as input to or produced as output from the programs of this Package do not automatically fall under the copyright of this Package, but belong to whomever generated them, and may be sold commercially, and may be aggregated with this Package. If such scripts or library files are aggregated with this Package via the so-called "undump" or "unexec" methods of producing a binary executable image, then distribution of such an image shall neither be construed as a distribution of this Package nor shall it fall under the restrictions of Paragraphs 3 and 4, provided that you do not represent such an executable image as a Standard Version of this Package.
- 8 C subroutines (or comparably compiled subroutines in other languages) supplied by you and linked into this Package in order to emulate subroutines and variables of the language defined by this Package shall not be considered part of this Package, but are the equivalent of input as in Paragraph 6, provided these subroutines do not change the language in any way that would cause it to fail the regression tests for the language.
- 9 Aggregation of this Package with a commercial distribution is always permitted provided that the use of this Package is embedded; that is, when no overt attempt is made to make this Package's interfaces visible to the end user of the commercial distribution. Such use shall not be construed as a distribution of this Package.
- 10 The name of the Copyright Holder may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS PACKAGE IS PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

The End

Pexpect

Copyright © 2010 Noah Spurrier

Credits: Noah Spurrier, Richard Holden, Marco Molteni, Kimberley Burchett, Robert Stone, Hartmut Goebel, Chad Schroeder, Erick Tryzelaar, Dave Kirby, Ids vander Molen, George Todd, Noel Taylor, Nicolas D. Cesar, Alexander Gattin, Geoffrey Marshall, Francisco Lourenco, Glen Mabey, Karthik Gurusamy, Fernando Perez, Corey Minyard, Jon Cohen, Guillaume Chazarain, Andrew Ryan, Nick Craig-Wood, Andrew Stone, Jorgen Grahn (Let me know if I forgot anyone.)

Free, open source, and all that good stuff.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

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 AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

PHP License

The PHP License, version 3.01

Copyright © 1999 - 2009 The PHP Group. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, is 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.
- 2 Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- 3 The name "PHP" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact group@php.net.
- 4 Products derived from this software may not be called "PHP", nor may "PHP" appear in their name, without prior written permission from group@php.net. You may indicate that your software works in conjunction with PHP by saying "Foo for PHP" instead of calling it "PHP Foo" or "phpfoo"
- 5 The PHP Group may publish revised and/or new versions of the license from time to time. Each version will be given a distinguishing version number.
 - Once covered code has been published under a particular version of the license, you may always continue to use it under the terms of that version. You may also choose to use such covered code under the terms of any subsequent version of the license published by the PHP Group. No one other than the PHP Group has the right to modify the terms applicable to covered code created under this License.
- 6 Redistributions of any form whatsoever must retain the following acknowledgment:
"This product includes PHP software, freely available from <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE PHP DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This software consists of voluntary contributions made by many individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project, please see <<http://www.php.net>>.

PHP includes the Zend Engine, freely available at <<http://www.zend.com>>.

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>

Robb Shecter <robb@acm.org>

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

Florent Rougon <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 is 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

The 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 Web site at

<http://www.vertica.com/licenses/pythondialog-2.7.tar.bz2>

<http://www.vertica.com/licenses/pythondialog-2.7.tar.bz2>

RRDTool Open Source License

Note: rrdtool is a dependency of using the ganglia-web third-party tool. RRDTool allows the graphs displayed by ganglia-web to be produced.

RRDTOOL - Round Robin Database Tool

A tool for fast logging of numerical data graphical display of this data.

Copyright © 1998-2008 Tobias Oetiker

All rights reserved.

GNU GPL License

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

This program 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 General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 59 Temple Place - Suite 330, Boston, MA 02111-1307, USA

FLOSS License Exception

(Adapted from <http://www.mysql.com/company/legal/licensing/foss-exception.html>)

I want specified Free/Libre and Open Source Software ("FLOSS") applications to be able to use specified GPL-licensed RRDtool libraries (the "Program") despite the fact that not all FLOSS licenses are compatible with version 2 of the GNU General Public License (the "GPL").

As a special exception to the terms and conditions of version 2.0 of the GPL:

You are free to distribute a Derivative Work that is formed entirely from the Program and one or more works (each, a "FLOSS Work") licensed under one or more of the licenses listed below, as long as:

- 1 You obey the GPL in all respects for the Program and the Derivative Work, except for identifiable sections of the Derivative Work which are not derived from the Program, and which can reasonably be considered independent and separate works in themselves
- 2 All identifiable sections of the Derivative Work which are not derived from the Program, and which can reasonably be considered independent and separate works in themselves
 - are distributed subject to one of the FLOSS licenses listed below, and
 - the object code or executable form of those sections are accompanied by the complete corresponding machine-readable source code for those sections on the same medium and under the same FLOSS license as the corresponding object code or executable forms of those sections.
- 3 Any works which are aggregated with the Program or with a Derivative Work on a volume of a storage or distribution medium in accordance with the GPL, can reasonably be considered independent and separate works in themselves which are not derivatives of either the Program, a Derivative Work or a FLOSS Work.

If the above conditions are not met, then the Program may only be copied, modified, distributed or used under the terms and conditions of the GPL.

FLOSS License List

License name	Version(s)/Copyright Date
Academic Free License	2.0
Apache Software License	1.0/1.1/2.0
Apple Public Source License	2.0
Artistic license	From Perl 5.8.0
BSD license	"July 22 1999"
Common Public License	1.0
GNU Library or "Lesser" General Public License (LGPL)	2.0/2.1
IBM Public License, Version	1.0
Jabber Open Source License	1.0
MIT License (As listed in file MIT-License.txt)	-
Mozilla Public License (MPL)	1.0/1.1
Open Software License	2.0
OpenSSL license (with original SSLeay license)	"2003" ("1998")
PHP License	3.0
Python license (CNRI Python License)	-
Python Software Foundation License	2.1.1
Sleepycat License	"1999"

W3C License	"2001"
X11 License	"2001"
Zlib/libpng License	-
Zope Public License	2.0/2.1

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 © 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:
- 6 "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>"
- 7 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, MERCHANTABILITY, 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) or our web site at (<http://www.spread.org/comments>).

SNMP

Various copyrights apply to this package, listed in various separate parts below. Please make sure that you read all the parts. Up until 2001, the project was based at UC Davis, and the first part covers all code written during this time. From 2001 onwards, the project has been based at SourceForge, and Networks Associates Technology, Inc hold the copyright on behalf of the wider Net-SNMP community, covering all derivative work done since then. An additional copyright section has been added as Part 3 below also under a BSD license for the work contributed by Cambridge Broadband Ltd. to the project since 2001. An additional copyright section has been added as Part 4 below also under a BSD license for the work contributed by Sun Microsystems, Inc. to the project since 2003.

Code has been contributed to this project by many people over the years it has been in development, and a full list of contributors can be found in the README file under the THANKS section.

Part 1: CMU/UCD copyright notice: (BSD like)

Copyright © 1989, 1991, 1992 by Carnegie Mellon University

Derivative Work - 1996, 1998-2000

Copyright © 1996, 1998-2000 The Regents of the University of California

All Rights Reserved

Permission to use, copy, modify and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appears in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of CMU and The Regents of the University of California not be used in advertising or publicity pertaining to distribution of the software without specific written permission.

CMU AND THE REGENTS OF THE UNIVERSITY OF CALIFORNIA DISCLAIM ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL CMU OR THE REGENTS OF THE UNIVERSITY OF CALIFORNIA BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM THE LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Part 2: Networks Associates Technology, Inc copyright notice (BSD)

Copyright © 2001-2003, Networks Associates Technology, Inc

All rights reserved.

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

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the Networks Associates Technology, Inc nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Part 3: Cambridge Broadband Ltd. copyright notice (BSD)

Portions of this code are copyright (c) 2001-2003, Cambridge Broadband Ltd.

All rights reserved.

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

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

- The name of Cambridge Broadband Ltd. may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDER ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Part 4: Sun Microsystems, Inc. copyright notice (BSD)

Copyright © 2003 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

Use is subject to license terms below.

This distribution may include materials developed by third parties.

Sun, Sun Microsystems, the Sun logo and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

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

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the Sun Microsystems, Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Part 5: Sparta, Inc copyright notice (BSD)

Copyright © 2003-2006, Sparta, Inc

All rights reserved.

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

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of Sparta, Inc nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Part 6: Cisco/BUPTNIC copyright notice (BSD)

Copyright © 2004, Cisco, Inc and Information Network Center of Beijing University of Posts and Telecommunications.

All rights reserved.

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

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of Cisco, Inc, Beijing University of Posts and Telecommunications, nor the names of their contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Part 7: Fabasoft R&D Software GmbH & Co KG copyright notice (BSD)

Copyright © Fabasoft R&D Software GmbH & Co KG, 2003

oss@fabasoft.com

Author: Bernhard Penz

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

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- The name of Fabasoft R&D Software GmbH & Co KG or any of its subsidiaries, brand or product names may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDER ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER BE

LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Tecla Command-line Editing

Copyright © 2000 by Martin C. Shepherd.

All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, provided that the above copyright notice(s) and this permission notice appear in all copies of the Software and that both the above copyright notice(s) and this permission notice appear in supporting documentation.

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 AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR HOLDERS INCLUDED IN THIS NOTICE BE LIABLE FOR ANY CLAIM, OR ANY SPECIAL INDIRECT OR CONSEQUENTIAL DAMAGES, OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Except as contained in this notice, the name of a copyright holder shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization of the copyright holder.

Webmin Open Source License

Copyright © Jamie Cameron

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.
- 2** Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- 3** Neither the name of the developer nor the names of contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE DEVELOPER "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE DEVELOPER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

xerces

NOTICE file corresponding to section 4(d) of the Apache License, Version 2.0, in this case for the Apache Xerces distribution.

This product includes software developed by The Apache Software Foundation (<http://www.apache.org/>).

Portions of this software were originally based on the following:

Software copyright © 1999, IBM Corporation., <http://www.ibm.com>.

zlib

This is used by the project to load zipped files directly by COPY command. www.zlib.net/

zlib.h -- interface of the 'zlib' general purpose compression library version 1.2.3, July 18th, 2005

Copyright © 1995-2005 Jean-loup Gailly and Mark Adler

This software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

- 1 The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.
- 2 Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
- 3 This notice may not be removed or altered from any source distribution.

Jean-loup Gailly jloup@gzip.org

Mark Adler madler@alumni.caltech.edu