
Vertica Database 2.0.5-0

Release Notes

Copyright© 2006, 2007 Vertica Systems, Inc.

Date of Publication: 1/8/2008



CONFIDENTIAL

Copyright Notice

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

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

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

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

Trademarks

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

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

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

Fedora™ is a trademark of Red Hat, Inc.

Intel® is a registered trademark of Intel.

Linux® is a registered trademark of Linus Torvalds.

Microsoft® is a registered trademark of Microsoft Corporation.

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

Oracle® is a registered trademark of Oracle Corporation.

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

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

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

Open Source Software Acknowledgements

Boost

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

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

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

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

PostgreSQL

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

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

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

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

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

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

Python Dialog

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

Upstream Author:

Peter Astrand <peter@cendio.se>

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

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

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

Spread

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

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

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

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

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

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

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

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

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

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

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

Contents

Copyright Notice	ii
<hr/>	
Technical Support	9
<hr/>	
About the Documentation	11
<hr/>	
Where to Find the Vertica Documentation	11
Reading the HTML Files	12
Printing the PDF Files.....	12
<hr/>	
Suggested Reading Paths	13
<hr/>	
Where to Find Additional Information	15
Typographical Conventions	16
<hr/>	
Vertica Database V2.0.5-0	17
<hr/>	
Supported Platforms	19
<hr/>	
Software Compatibility Issues	20
<hr/>	
Version Compatibility Issues	21
<hr/>	
New and Changed Features	22
<hr/>	
Advanced Features.....	22
Database Administration.....	23
Database Designer	24
Installation	25
Licensing	25
Performance.....	26
SQL Data Definition Language	27
SQL Data Manipulation Language	27
SQL Query Syntax.....	29
vsqL.....	30
<hr/>	
Known Problems and Workarounds	31
<hr/>	

Release Notes

Administration Tools	31
Catalog and Data Pathname Restrictions	31
Database Name Restriction	31
Exiting Catalog and Data Directories Cause Database Creation to Fail	31
Installation	31
Hostname Resolution May Cause Installation to Fail	31
Installation May Fail to Set Ownership of Catalog and Data Directories	31
Installation May Fail to Set Up Passwordless SSH	32
Invalid RPM Package Parameter Causes Installation Script Failure	32
Old AMD Opteron Chips Not Supported	32
PermitRootLogin Must Be Enabled	33
Reinstallation Requires Same Host List	33
NFS Mounted or Symbolically Linked Directories Not Supported For Catalog/Data	33
SQL	33
ALTER TABLE Restrictions	33
COMMIT May Fail Under Heavy Server Load	33
COPY From Standard Input Handles NULL Symbol Incorrectly	34
Inconsistent Transaction Rollback	34
INSERT SELECT With No FROM Clause Fails	34
System Administration	35
Create Database Requires Mount Point /dev/pts	35
Database Hangs Under Rare Conditions	35
Database Startup Requires All Nodes Available	35
Refresh May Prevent Database Shutdown	35

Problems Fixed 37

Documentation Notes 39

Vertica V1.3.16-0 41

New and Changed Features	41
Problems Fixed	41
Known Problems and Workarounds	42
Documentation Notes	43

Index 45

Technical Support

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

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

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

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

About the Documentation

Where to Find the Vertica Documentation

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

Database Server Systems

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

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

Database Client Systems

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

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

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

World Wide Web

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

http://www.vertica.com/v-zone/product_documentation `http://www.vertica.com/v-zone/product_documentation`.

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

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

Reading the HTML Files

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

The Vertica documentation has been tested on the following browsers:

- Internet Explorer 7
- FireFox
- Opera
- Safari

Please report any script, image rendering, or text formatting problems to **Technical Support** (on page 9).

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

Printing the PDF Files

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

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

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

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

Suggested Reading Paths

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

All Users

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

System Administrators

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

Database Administrators

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

Application Developers

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

Where to Find Additional Information

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

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

Typographical Conventions

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

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

Vertica Database V2.0.5-0

Vendor Information

Vertica Systems, Inc.

Three Dundee Park Drive, Suite 102

Andover, MA 01810-3723

Phone: (978) 475-1070

Fax: (978) 475-6855

E-Mail: info@vertica.com

Web site: <http://www.vertica.com> (<http://www.vertica.com>)

Product and Version

Vertica Database V2.0.5-0

Media

The Vertica software and documentation are available for download (with registration) on the **Vertica website download page** http://www.vertica.com/v-zone/download_vertica. To obtain these materials on DVD, contact **Technical Support** (on page 9).

Supported Platforms

Hardware

See the **Vertica Web site** (<http://www.vertica.com>) for hardware requirements and recommendations.

Operating Systems for Vertica Database Servers

Supported for all purposes including production:

- Red Hat Enterprise Linux 4, 64 bit
- Red Hat Enterprise Linux 5, 64 bit
- SUSE Linux Enterprise Server 10, 64 bit

Supported for non-production (testing and evaluation) purposes only:

- Red Hat Enterprise Linux 4, 32 bit
- Red Hat Enterprise Linux 5, 32 bit
- SUSE Linux Enterprise Server 10, 32 bit
- Fedora Core 6, 32 bit
- Fedora Core 6, 64 bit

Not Supported

- VMware

You can install Vertica on VMware virtual machines running Red Hat Enterprise Linux 4, 32 bit or 64 bit. However, Vertica Systems, Inc. strongly suggests that you **do not use this platform for any sort of storage capacity or performance testing.**

Operating Systems for Vertica Client Applications

Any operating system that has the following items installed:

- Java Virtual Machine (JVM)
- Vertica JDBC jar file included in the application's CLASSPATH environment variable

Software Compatibility Issues

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

- remove or disable all non-essential applications from cluster hosts
- deploy Vertica as the only active process on each host (other than Linux processes)

Version Compatibility Issues

- Vertica **V1.0** and earlier must be **uninstalled**. Databases built with **V1.0** are not compatible with newer versions of Vertica. You must **rebuild** the database using a later version.
- Vertica **V1.1** can be upgraded to **V1.2**. Use the instructions in the V1.2 1-9 documentation.
- Vertica **V1.2** can be upgraded to **V1.3**. Use the instructions in the V1.3.16-0 documentation.
- Vertica **V1.3** can be upgraded to **2.0**. Follow the instructions in the Installation Guide.

New and Changed Features

This section describes features have been introduced or changed in Vertica V2.0.3-0.

Advanced Features

Adding a New Projection

The CREATE PROJECTION command (described in the SQL Reference Manual) can now be used to add a projection to a database when one or more of the tables named in the projection query already contains data. However, a newly-created projection is not available for query processing until you execute a refresh, which copies data into the projection.

Make sure that the new projection has a sufficient number of buddy projections and that all nodes are up before using the new SQL START_REFRESH function to initiate a refresh. You can check the progress of the refresh by monitoring the log files or calling the following SQL functions (described in the SQL Reference Manual):

- GET_PROJECTION_STATUS
- GET_TABLE_PROJECTIONS

Automatic Replication of Dimension Tables

The CREATE PROJECTION command (described in the SQL Reference Manual) now supports the automatic replication of small dimension table projections using the following syntax:

```
UNSEGMENTED { NODE node | ALL NODES }
```

where *node* allows you to replicate the projection only on a specific node.

For example, if the nodes are named NODE01, NODE02, and NODE03 then:

```
CREATE PROJECTION ABC ... UNSEGMENTED ALL NODES
```

creates projections named ABC_NODE01, ABC_NODE02, and ABC_NODE03.

The replication of dimension tables is required for K-Safety.

Hash Segmentation

The CREATE PROJECTION command (described in the SQL Reference Manual) now supports hash segmentation using the following syntax:

```
SEGMENTED BY HASH ( column [ , ... ] )  
[ ALL NODES [ OFFSET offset ] | NODES node [ , ... ] ]
```

where *column* is a table column and OFFSET allows you to create a buddy projection (required for K-Safety) using a different node on which to start the segmentation distribution.

Range segmentation is still supported and both forms of segmentation allow the keyword `SITE` (in place of `NODE`) for backwards compatibility.

For More Information

See the following sections of the Database Administrator's Guide (Advanced):

- Adding Custom Projections
- Automatic Replication of Unsegmented Projections
- Hash Segmentation
- Range Segmentation

Database Administration

Administration Tools

The Administration Tools now include:

- online help in the form of a Main Menu command and Help buttons. See Help on Using the Administration Tools in the Database Administrator's Guide.
- the ability to control debug logging. See Database Parameters in the Database Administrator's Guide.
- an informative error message that appears when attempting to shut down a database that still has connections open:
`NOTICE: Cannot shut down while users are connected`
See Stop Database in the Database Administrator's Guide for more information.

Logging and Diagnostics

- The log files `dbLog` and `vertica.log` have been consolidated into a single file (`vertica.log`) that contains all messages issued when a database is running.
- `dbLog` continues to exist but is now only used to capture messages when no database is running.
See Monitoring the Log Files in the Troubleshooting Guide.

Monitoring API

Vertica provides an API for monitoring various features and functions within a database in the form of virtual tables that can be queried using a limited form of the `SELECT` statement. You can use external tools to query the virtual tables and act upon the information as desired. For example, you can program a third-party monitoring tool to periodically query the K-Safety level of the database.

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

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

Database Designer

User Interface

The Database Designer user interface has been redesigned for ease of use. It now allows you to specify:

- the location in which to write the design (SQL projection script) and a log file containing information about the design produced.
- whether to produce a basic design (single-table superprojections only) or an optimized design (superprojections and pre-join projections).

Functional Dependencies

The Database Designer now produces physical schema designs that take functional dependencies into consideration for increased query performance. A functional dependency is a relationship between two sets of column values in a table where one column value can determine the other. It represents knowledge of the data that cannot otherwise be expressed in the logical schema. In Vertica functional dependencies are expressed as `CORRELATION` constraints and are used by the Database Designer to produce more efficient physical schema designs.

Use the `ALTER TABLE` command to add `CORRELATION` constraints to a table, as described in the SQL Reference Manual.

For More Information

In the Database Administrator's Guide see:

- Using the Database Designer
- Create New Database Design
- Database Designer Dialog
- Modify Existing Database Design

Installation

Vertica now supports the use of IP addresses in place of hostnames in the installation script. Make sure not include space characters in the list. For example:

```
# /opt/vertica/sbin/install_vertica -s 192.168.233.101,192.168.233.102,192.168.233.103 ...
```

See Installing Vertica in the Installation Guide for more information.

Licensing

Vertica now supports license keys that provide full product functionality for a specific period of time or forever if you purchase a perpetual license. There is a grace period during which Vertica continues to work after the license has expired. The length of the grace period depends on the license.


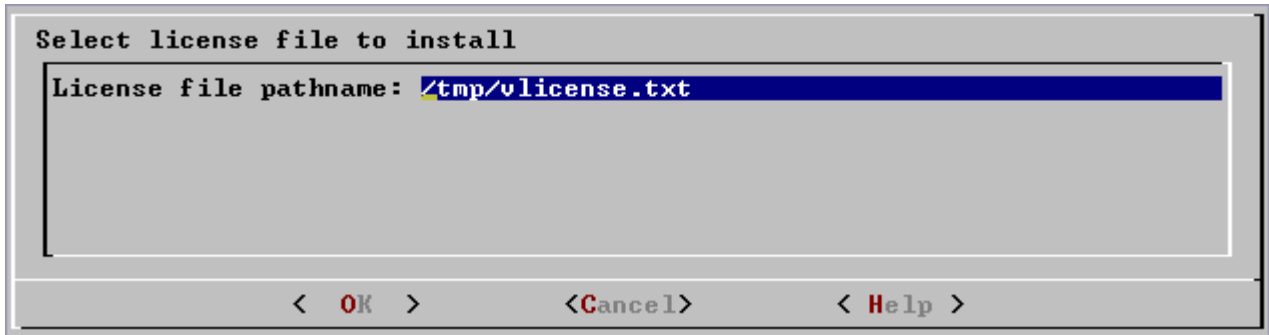
Obtaining a License Key

1. Go to <http://www.vertica.com/support> (<http://www.vertica.com/support>) and follow the instructions there. You will be given download access to a file whose name includes your organization and the date the license was requested.
2. Download the license file to the `/tmp` directory on the Administration Host.

Installing a New or Upgrade Installation License Key

1. Obtain a license key as described above.
2. Install Vertica as described in the Installation Guide.

3. The first time you log in as the Database Administrator and run the Vertica Administration Tools, the user interface displays a form asking for the path to the license key file. Enter the **absolute pathname** of the file of the form and press OK.



For more information, see Using the Administration Tools in the Database Administrator's Guide.

Installing a Renewal License Key

1. Obtain a license key as described above.
2. Start a database. The database can run with an expired license but you cannot select data.
3. In the Administration Tools, go to Advanced > Upgrade License Key.
4. Select the database for which to upgrade the license key.
5. Enter the absolute pathname of your downloaded license key file and click OK.

For more information, see Upgrade License Key in the Database Administrator's Guide.

Using Vertica Functions to Manage License Keys

The following functions are described in the SQL Reference Manual:

- Use the `INSTALL_LICENSE` function to install the license key in the global catalog.
- Use the `DISPLAY_LICENSE` function to display the license information.

Performance

Several enhancements have been made to increase query performance, all of which are transparent to users of Vertica.

SQL Data Definition Language

ALTER TABLE

The ALTER TABLE command (described in the SQL Reference Manual) now includes:

- the **ADD COLUMN** clause, which automatically adds a new column to a table and to all superprojections of the table. The encoding type for the new column is set to AUTO.
- **CORRELATION** table constraints, which capture functional dependencies used by the *Database Designer* (page 24) to generate optimized projections.

DROP TABLE

The DROP TABLE command (described in the SQL Reference Manual) now requires the CASCADE keyword, which drops all projections that include the table. Only the syntax of the command has changed. The behavior of the command is the same as in previous versions of Vertica.

SQL Data Manipulation Language

COPY

The COPY command (described in the SQL Reference Manual) now:

- loads specific columns
- allows a user-defined line terminator
- logs exceptions and writes rejected rows to files
- can abort on error and roll back

DELETE/UPDATE

The DELETE and UPDATE commands (described in the SQL Reference Manual) are now enabled by default. The behavior of the commands has not changed:

- DELETE and UPDATE always write tuples to the WOS.
- UPDATE writes two tuples into the WOS: one with new data and one marked for deletion.
- DELETE does not actually delete data from disk storage; it simply marks tuples as deleted so that they can be found by historical queries. The ability to delete data from disk storage will be available in a future release.

LCOPY

The new LCOPY command (described in the SQL Reference Manual) is identical to the COPY command except that it loads data from a client system, rather than a cluster host.

Large Result Set Support in ODBC

Three new attributes have been added for large result set support. They are defined in the file `verticaodbc.h`.

1. The **connection attribute ATTR_VERTICA_LRS_PATH** specifies the location in which the ODBC driver keeps temporary files for large result sets. For example:

```
char * lrs_path="/my_disk/tmp";
ret = SQLSetConnectAttr(conn.dbc, SQL_ATTR_VERTICA_LRS_PATH,
    (PTR)lrs_path, strlen(lrs_path));
```

Unix/Linux default values:

- If the environment variable TMPDIR exists and contains the name of an appropriate directory, that is used.
- Otherwise, if the dir argument is non-NULL and appropriate, it is used.
- Otherwise, "/tmp" is used.

Windows default values:

- If the TMP environment variable is defined and set to a valid directory name, that name is used.
- Otherwise, the dir parameter is used as the path.
- If the dir parameter is NULL or set to the name of a directory that does not exist, the current working directory is used.

2. The **statement attribute SQL_ATTR_VERTICA_MAX_MEM_CACHE** defines the maximum memory for the client storage of a large result set. If the result set size exceeds this value, the ODBC driver uses a temporary file to keep the large result set or uses streaming mode for fetching data from the database server. For example:

```
SQLINTEGER mem_cache_size=256*1024*1024; // 256 MB
SQLSetStmtAttr(hstmt, SQL_ATTR_VERTICA_MAX_MEM_CACHE,
    (PTR)mem_cache_size, 0);
```

The default value is 64MB.

3. The **statement attribute SQL_ATTR_VERTICA_LRS_STREAMING** specifies that the ODBC driver should use a temporary file to keep the large result set, or use streaming mode to fetch the large result set from the database server. If the value is TRUE, the ODBC driver pauses the query execution when the memory cache on the client is full and resumes execution of the query after the memory cache rows are retrieved by the ODBC application using `SQLFetch`. For example:

```
SQLINTEGER lrs_streaming=1;
SQLSetStmtAttr(hstmt, SQL_ATTR_VERTICA_LRS_STREAMING,
    (PTR)lrs_streaming, 0);
```

The default value is FALSE.

SQL Query Syntax

New Aggregate Functions

Supported aggregate functions now include:

- STDDEV
- STDDEV_POP
- STDDEV_SAMP
- VAR_POP
- VAR_SAMP
- VARIANCE

See Aggregate Functions in the SQL Reference Manual.

DISTINCT Aggregates

The DISTINCT set-quantifier keyword can now be used before the argument to the aggregate functions COUNT, MIN, MAX, AVG, and SUM. For example:

```
SELECT COUNT (DISTINCT customer_key)
FROM fact
```

See Aggregate Expressions in the SQL Reference Manual.

Expressions in GROUP BY

The GROUP BY clause of a query can now contain expressions. For example:

```
SELECT  c1+c2, c2+32, COUNT(*)
FROM    fact
GROUP BY c1+c2, c2+32;
```

See the GROUP BY Clause in the SQL Reference Manual.

Compound Predicate Joins

Tables can now be joined based on compound (multiple-column) primary and foreign keys. For example:

```
SELECT *
FROM fact, dimension
WHERE fact.fk1 = dimension.pk1
      AND
      fact.fk2 = dimension.pk2;
```

See Using Compound Predicate Joins in the SQL Programmer's Guide.

Multiple Dimension Table References in the WHERE Clause

The same dimension table can now appear in multiple predicates in the WHERE clause of a query. For example:

```
SELECT *
FROM fact, dimension
WHERE fact.fk = dimension.pk
      AND
      fact.name = dimension.name
```

See Using Multiple Dimension Table References in the WHERE Clause in the SQL Programmer's Guide.

Multiple Instances of Dimension Tables in the FROM Clause

The same dimension table can now appear multiple times in the FROM clause of a query, using different aliases. For example:

```
SELECT *
FROM fact, dimension d1, dimension d2
WHERE fact.fk = d1.pk
      AND
      fact.name = d2.name
```

See Using Multiple Instances of Dimension Tables in the FROM Clause in the SQL Programmer's Guide.

vsqI

The `-w password` command line option now allows you to specify a database user password on the command line.

Using this command line option displays the database password in plain text on the screen. Use it with care particularly if you are connecting as the Database Administrator.

Known Problems and Workarounds

The following are known problems in Vertica V2.0.3-0.

Administration Tools

Catalog and Data Pathname Restrictions

Failure to comply with these restrictions may result in database creation failure:

- Catalog and data pathnames must contain only alphanumeric characters.
- Catalog and data pathnames cannot have leading space characters.

Database Name Restriction

Database names must begin with an alphabetic character.

Exiting Catalog and Data Directories Cause Database Creation to Fail

Creating a new database fails when the specified catalog and/or data directories already exist on at least one host. To correct the problem, manually remove those directories and create the database again.

Installation

Hostname Resolution May Cause Installation to Fail

The following error indicates that the hostname specified for installation does not match the hostname returned by the system.

```
Hostname resolution (FAILED)
```

See Check Hostname Resolution in the Installation Guide for instructions.

Installation May Fail to Set Ownership of Catalog and Data Directories

The `install_vertica` script described in the `install_vertica` section of the Installation Guide allows you to provide the optional parameter `data_directory`, which specifies the default directory for database data and catalog files. If you specify a `data_directory` that is owned by a user that no longer exists, `install_vertica` silently fails to set the correct permissions on the directory. Creating a database using that directory fails as a result.

Workaround

As root, use chown to set the ownership of the directory to the Database Administrator.

Installation May Fail to Set Up Passwordless SSH

If the installation process fails to establish passwordless ssh access to all hosts for the Database Administrator, see the Enable Secure Shell (SSH) Logins section of the Installation Guide.

Invalid RPM Package Parameter Causes Installation Script Failure

If you specify an invalid RPM package name as a parameter for the install_vertica script, the installation may fail with the following error:

```
Error: need more than 0 values to unpack
Installation failed.
```

Old AMD Opteron Chips Not Supported

Problem

Certain operations involving dividing timestamps or date ranges can cause the server to fail.

Explanation

Very old versions of AMD Opteron chips are not supported. These chips do not implement the SSE3 instruction set extensions, which the Vertica server uses.

Workaround

Use the following steps to determine whether or not your processor is affected:

1. Determine your processor model. In Linux run the command:

```
cat /proc/cpuinfo | grep "model name"
```

For example:

```
model name : AMD Opteron(tm) Processor 246
model name : Intel(R) Pentium(R) M processor 2.00GHz
```

2. If your machine has an AMD Opteron processor with a three digit model number where the middle digit is 4, your processor may be affected if it was manufactured before August, 2005. If the processor is not supported, its part number is listed under one of the "Sledgehammer" headings on the following page:

http://en.wikipedia.org/wiki/List_of_AMD_Opteron_microprocessor
http://en.wikipedia.org/wiki/List_of_AMD_Opteron_microprocessor

You may find this confusing because AMD used the same model name for different part numbers. For example, an "Opteron 246" might be an affected OSA246CEP5AL or it might be an unaffected OSA246FAA5BL.

PermitRootLogin Must Be Enabled

The installation process requires that the entry for `PermitRootLogin` in `/etc/ssh/sshd_config` be set to `Yes`. Otherwise you may get the following failure message:

```
Error: 1
Traceback (most recent call last):
  File "/opt/vertica/bin/verticaInstall.py", line 816, in ?
    DBfunctions.record("Error: "+str(e)+"\n"+traceback.format_exc())
AttributeError: 'module' object has no attribute 'format_exc'
Installation failed.
```

Reinstallation Requires Same Host List

If an installation fails and you attempt the installation again, provide the same list of host names that you did the first time. For example, if you entered:

```
# /opt/vertica/sbin/install_vertica -s host01,host02,host03 -r ...
```

and the installation fails, use the same command the second time. If you use a different command, the installation will fail again and you will have to clean up the resulting files manually. Contact **Technical Support** (on page 9) for assistance.

NFS Mounted or Symbolically Linked Directories Not Supported For Catalog/Data

Vertica installation does not support nfs mounted or symbolically linked catalog or data directories.

SQL

ALTER TABLE Restrictions

- You cannot add NOT NULL or any other column constraints to a table that already has projections defined. Define column constraints at CREATE TABLE time.
- You can use only one ADD COLUMN clause at a time.
- Use multiple ALTER TABLE statements to add multiple columns.
- You cannot add a column to a temporary table. An attempt to do so causes the database to crash.

COMMIT May Fail Under Heavy Server Load

Under rare circumstances, a COMMIT command may fail due to heavy database server load. Poor server performance may serve as a warning.

COPY From Standard Input Handles NULL Symbol Incorrectly

The COPY command loses one of the backslashes from the NULL string on the command line. For example:

```
cat test | Linux/bin/vsql -a -c " copy a from stdin delimiter ',' null '\\N';"
```

The server receives a NULL string of "N" rather than "\N" as was specified in the command. The workaround is to provide a third backslash to the command:

```
cat test | Linux/bin/vsql -a -c " copy a from stdin delimiter ',' null '\\\\N';"
```

Inconsistent Transaction Rollback

SQL parser and syntax errors do not cause a transaction rollback but other SQL semantic or query execution errors do. For example:

```
=> SELECT COUNT(*) FROM T;
count
-----
13
(1 row)

=> INSERT INTO T VALUES(1,1);
OUTPUT
-----
1
(1 row)

=> INSERT INTO P VALUES(2,2);
ERROR: p is not a table. DML not supported
=> SELECT COUNT(*) FROM T;
count
-----
14
(1 row)

=> SELECT * FROM T T1, T T2 WHERE T1.C0 = T1.C1;
ROLLBACK: Not a Star or Snow-Flake Query
=> SELECT COUNT(*) FROM T;
count
-----
13
(1 row)
```

INSERT SELECT With No FROM Clause Fails

The following query fails:

```
INSERT INTO t SELECT 'aaa', 'bbb', 10;
```

Use INSERT ... VALUES when no column values contain column references:

```
INSERT INTO t VALUES 'aaa','bbb',10;
```

System Administration

Create Database Requires Mount Point /dev/pts

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

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

Before you install, make sure that /dev/pts is mounted.

Database Hangs Under Rare Conditions

The system may appear to hang under the following rare conditions

- the system is overloaded and/or running under resource constraints
- the first query in a user session results in an error due to resource constraints
- many connections are continuously executing SQL statements resulting in errors that force a rollback
- some connections are executing SQL statements resulting in errors and a connection is interrupted (Ctrl-C).

To work around these issues, disconnect all users and use the Administration Tools to stop and restart the database. If unsuccessful, use the Stop Vertica on Node (or if necessary Kill Vertica Process on Node) tool to stop all nodes, then restart the database.

Database Startup Requires All Nodes Available

All nodes in a cluster must be available when starting a Database. If one node is not accessible, the database cannot be started on the cluster. If you encounter this problem, contact **Technical Support** (on page 9).

Refresh May Prevent Database Shutdown

Under certain circumstances, a refresh operation may prevent a database shutdown from completing. If you experience this problem, run diagnostics and contact **Technical Support** (on page 9) immediately.

Problems Fixed

Invalid Hostnames Containing Periods

Hostnames containing the period character were not handled correctly at database creation time. The problem has been corrected.

Create Database Failed But Appeared to Succeed

After creating a new database, the name did not appear in the list of databases to start. The create operation failed but the error messages could not be seen. The problem has been corrected.

Shutdown Error Message Confusing

Previously, open client connections were not correctly reported at shutdown time. The problem has been corrected.

INSERT ... SELECT Without FROM Clause Crashed Database

INSERT ... SELECT with a select list consisting only of constant values and no FROM clause should behave as if it were INSERT ... VALUES but caused the database to crash. The problem has been corrected.

INSERT /*+ direct */ SELECT with Expression and ORDER BY Crashed Database

An INSERT statement that had all the following items caused the database to crash:

- a /*direct*/ hint
- an expression in the select list
- an ORDER BY clause

The problem has been corrected.

Lock Timeout Error Due to SERIALIZABLE Isolation Not Helpful

Previously, a lock timeout error message did not indicate the source of the problem. The error message now indicates that the problem is a result of the SERIALIZABLE isolation level.

Connecting With Password Crashes Databases

Invoking vsql with the `-w password` option caused the database to crash. The problem has been corrected.

JDBC Access to Catalog Caused Exception

The JDBC describe and select from catalog operations failed with a null pointer exception. They now return an error message.

ODBC Driver Returned Incorrect Null Values

Queries executed using the ODBC driver return incorrect null values when data was present. The problem has been corrected.

JDBC Request for Isolation Level Failed

The JDBC function `Connection.getTransactionIsolation()` returned an error. The problem has been corrected.

Database Designer Failed When Using \t as Delimiter

Running the Database Designer with a sample data column delimiter of \t caused an incorrect column count. The problem has been corrected.

Database Designer Failed When Trying to Load Statistics

The Database Designer failed to produce a projection script due to a failure to load table and projection statistics. The problem has been corrected.

Documentation Notes

The following notes apply to the Vertica V2.0.5-0 documentation:

Installation Guide

The description of the installation procedure for single-node installations has been clarified.

Release Notes

The Release Notes have been somewhat reorganized with respect to information from previous releases.

SQL Reference Manual

The SQL Reference Manual has been reorganized and new sections have been added.

Vertica V1.3.16-0

New and Changed Features

The following features have been introduced or changed in Vertica V1.3.16-0:

Administration Tools

- The **Configure Nodes** dialog has been removed. Node definitions are now created automatically as part of Create Database. Installing and uninstalling the database software on a cluster is now done by the new installation script (see below).
- The **Configure Databases** menu has been removed. Create Database, Drop Database, and View Database are now under the **Configuration Menu**.
- The **Database Designer** is now under the **Configuration Menu**.
- In the **Advanced** menu:
 - **Roll Back Database to Catalog Version** has been changed to **Restart Database With Catalog Version**.
 - **Check Spread** has been removed. Spread is now configured at installation time.

Database Designer

The Database Designer now prints status messages to the standard output while it is working.

Installation and Configuration

Installing and configuring Vertica is now much easier. Almost all of the manual work is now done by a new shell script. The procedure is described in detail in the Installation Guide. The manual procedures are still provided in case of problems.

Load Performance

Load performance has improved due to fewer locks held by the Tuple Mover.

Tuple Mover

The Tuple Mover is now more efficient and performs mergeouts in less time.

Problems Fixed

The following problems were fixed in Vertica V1.3.16-0:

Administration Tools "Install/Upgrade DB software on node"

Previously, the Administration Tools crashed when performing "Install/Upgrade DB software on node" operations due to certain older builds that could not be upgraded. The Administration Tools now produce an error message indicating that the currently-installed build requires manual uninstallation and reinstallation.

Database Designer

The Database Designer now produces an error message when the `sample data file directory:` field is empty (the default value is deleted and not replaced).

Diagnostics

`ERROR unrecognized time zone name: ""` problem has been fixed.

SQL and Query Execution

- Null constraints are now enforced when using the INSERT statement.
- An `INSERT /*+ DIRECT */ ... (SELECT ...)` statement's query failed when the SELECT list contained an expression and an ORDER BY clause was specified. The problem has been fixed.
- `COPY ... DIRECT` failed with VAssert when using multiple load streams. The problem has been fixed.

Known Problems and Workarounds

The following are known problems in Vertica V1.3.16-0:

Network Communication Test Fails

You may see the following FAILED message in the output from the `install_vertica` script:

```
root@hostname: Network communication (FAILED)
root@hostname: -----
root@hostname: Site hostname did not receive message from hostname
```

Explanation

This error occurs when hostnames are longer than 22 characters.

Workaround

On a single-node cluster, ignore the FAILED message.

On a multiple-node cluster, edit the `/etc/hosts` file on each host in the cluster and add a shorter alias to each host's IP address. For example, if using fully qualified domain names, the `/etc/hosts` file on each host should contain something like this:

```
# Do not remove the following line, or various programs
# that require network functionality will fail.
127.0.0.1 localhost.localdomain localhost
192.168.233.101 host01.fully.qualified.com host01
192.168.233.102 host02.fully.qualified.com host02
192.168.233.103 host03.fully.qualified.com host03
```

Documentation Notes

The following notes apply to the Vertica V1.3.16-0 documentation:

Browser Compatibility

- All problems related to Internet Explorer have been fixed:
 - The Search tab now displays and works correctly.
 - Images now display correctly.
 - Nested lists now display correctly.
- The documentation will now be tested against Safari in addition to Internet Explorer, FireFox, and Opera.

Installation Guide

- The scope of the installation procedure documented in this manual has expanded to include steps that were previously described in the Database Administrator's Guide, specifically those related to creating nodes.
- Many steps that had to be done manually in previous versions are now done automatically.

Quick Start

- The single-node and multiple-node tutorial procedures have been combined into a single procedure. The steps are the same regardless of the number of nodes in the cluster.

Index

A

- About the Documentation • 9
- Acrobat • 13
- Administration Tools • 29
- Adobe Acrobat • 13
- Advanced Features • 20
- ALTER TABLE Restrictions • 31

B

- Backslash • 14
- Bold text • 14
- Braces • 14
- Brackets • 14

C

- Catalog and Data Pathname Restrictions • 29
- Colored bold text • 14
- COMMIT May Fail Under Heavy Server Load • 32
- COPY From Standard Input Handles NULL Symbol Incorrectly • 32
- Copyright Notice • ii
- Create Database Requires Mount Point /dev/pts • 33

D

- Database Administration • 21
- Database Designer • 22, 25
- Database Hangs Under Rare Conditions • 33
- Database Name Restriction • 29
- Database Startup Requires All Nodes Available • 33
- Documentation • 13
- Documentation Notes • 37, 41

E

- Ellipses • 14
- Exiting Catalog and Data Directories Cause Database Creation to Fail • 29

H

- Hostname Resolution May Cause Installation to Fail • 29
- HTML • 13

I

- Inconsistent Transaction Rollback • 32
- Indentation • 14
- INSERT SELECT With No FROM Clause Fails • 33
- Installation • 23, 29
- Installation May Fail to Set Ownership of Catalog and Data Directories • 30
- Installation May Fail to Set Up Passwordless SSH • 30
- Invalid RPM Package Parameter Causes Installation Script Failure • 30
- Italic text • 14

K

- Known Problems and Workarounds • 29, 40

L

- Licensing • 23

M

- Monospace text • 14

N

- New and Changed Features • 20, 39
- NFS Mounted or Symbolically Linked Directories Not Supported For Catalog/Data • 31

O

- Old AMD Opteron Chips Not Supported • 30

P

- PDF • 13
- Performance • 24
- PermitRootLogin Must Be Enabled • 31
- Printing the PDF Files • 10
- Problems Fixed • 35, 39

R

- Reading the HTML Files • 10
- Refresh May Prevent Database Shutdown • 34
- Reinstallation Requires Same Host List • 31

S

- Shell script • 14
- Software Compatibility Issues • 18
- SQL • 31
- SQL Data Definition Language • 25

SQL Data Manipulation Language • 25

SQL Query Syntax • 27

Suggested Reading Paths • 11

Support • 7

Supported Platforms • 17

Syntax conventions • 14

System Administration • 33

T

Technical Support • 7, 10, 15, 31, 33, 34

Typographical Conventions • 14

U

Uppercase text • 14

V

Version Compatibility Issues • 19

Vertica Database V2.0.5-0 • 15

Vertica V1.3.16-0 • 39

Vertical line • 14

vsq1 • 28

W

Where to Find Additional Information • 13

Where to Find the Vertica Documentation • 9