



Informatica® PowerExchange for Vertica
10.1.1

User Guide for PowerCenter

© Copyright Informatica LLC 2014, 2018

This software and documentation are provided only under a separate license agreement containing restrictions on use and disclosure. No part of this document may be reproduced or transmitted in any form, by any means (electronic, photocopying, recording or otherwise) without prior consent of Informatica LLC.

Informatica, the Informatica logo, PowerCenter, and PowerExchange are trademarks or registered trademarks of Informatica LLC in the United States and many jurisdictions throughout the world. A current list of Informatica trademarks is available on the web at <https://www.informatica.com/trademarks.html>. Other company and product names may be trade names or trademarks of their respective owners.

Portions of this software and/or documentation are subject to copyright held by third parties, including without limitation: Copyright DataDirect Technologies. All rights reserved. Copyright © Sun Microsystems. All rights reserved. Copyright © RSA Security Inc. All Rights Reserved. Copyright © Ordinal Technology Corp. All rights reserved. Copyright © Aandacht c.v. All rights reserved. Copyright Genivia, Inc. All rights reserved. Copyright Isomorphic Software. All rights reserved. Copyright © Meta Integration Technology, Inc. All rights reserved. Copyright © Intalio. All rights reserved. Copyright © Oracle. All rights reserved. Copyright © Adobe Systems Incorporated. All rights reserved. Copyright © DataArt, Inc. All rights reserved. Copyright © ComponentSource. All rights reserved. Copyright © Microsoft Corporation. All rights reserved. Copyright © Rogue Wave Software, Inc. All rights reserved. Copyright © Teradata Corporation. All rights reserved. Copyright © Yahoo! Inc. All rights reserved. Copyright © Glyph & Cog, LLC. All rights reserved. Copyright © Thinkmap, Inc. All rights reserved. Copyright © Clearpace Software Limited. All rights reserved. Copyright © Information Builders, Inc. All rights reserved. Copyright © OSS Nokalva, Inc. All rights reserved. Copyright Edifecs, Inc. All rights reserved. Copyright Cleo Communications, Inc. All rights reserved. Copyright © International Organization for Standardization 1986. All rights reserved. Copyright © ej-technologies GmbH. All rights reserved. Copyright © Jaspersoft Corporation. All rights reserved. Copyright © International Business Machines Corporation. All rights reserved. Copyright © yWorks GmbH. All rights reserved. Copyright © Lucent Technologies. All rights reserved. Copyright © University of Toronto. All rights reserved. Copyright © Daniel Veillard. All rights reserved. Copyright © Unicode, Inc. Copyright IBM Corp. All rights reserved. Copyright © MicroQuill Software Publishing, Inc. All rights reserved. Copyright © PassMark Software Pty Ltd. All rights reserved. Copyright © LogiXML, Inc. All rights reserved. Copyright © 2003-2010 Lorenzi Davide, All rights reserved. Copyright © Red Hat, Inc. All rights reserved. Copyright © The Board of Trustees of the Leland Stanford Junior University. All rights reserved. Copyright © EMC Corporation. All rights reserved. Copyright © Flexera Software. All rights reserved. Copyright © Jinfonet Software. All rights reserved. Copyright © Apple Inc. All rights reserved. Copyright © Telerik Inc. All rights reserved. Copyright © BEA Systems. All rights reserved. Copyright © PDFlib GmbH. All rights reserved. Copyright © Orientation in Objects GmbH. All rights reserved. Copyright © Tanuki Software, Ltd. All rights reserved. Copyright © Ricebridge. All rights reserved. Copyright © Sencha, Inc. All rights reserved. Copyright © Scalable Systems, Inc. All rights reserved. Copyright © jqWidgets. All rights reserved. Copyright © Tableau Software, Inc. All rights reserved. Copyright © MaxMind, Inc. All Rights Reserved. Copyright © TMate Software s.r.o. All rights reserved. Copyright © MapR Technologies Inc. All rights reserved. Copyright © Amazon Corporate LLC. All rights reserved. Copyright © Highsoft. All rights reserved. Copyright © Python Software Foundation. All rights reserved. Copyright © BeOpen.com. All rights reserved. Copyright © CNRI. All rights reserved.

This product includes software developed by the Apache Software Foundation (<http://www.apache.org/>), and/or other software which is licensed under various versions of the Apache License (the "License"). You may obtain a copy of these Licenses at <http://www.apache.org/licenses/>. Unless required by applicable law or agreed to in writing, software distributed under these Licenses is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the Licenses for the specific language governing permissions and limitations under the Licenses.

This product includes software which was developed by Mozilla (<http://www.mozilla.org/>), software copyright The JBoss Group, LLC, all rights reserved; software copyright © 1999-2006 by Bruno Lowagie and Paulo Soares and other software which is licensed under various versions of the GNU Lesser General Public License Agreement, which may be found at <http://www.gnu.org/licenses/lgpl.html>. The materials are provided free of charge by Informatica, "as-is", without warranty of any kind, either express or implied, including but not limited to the implied warranties of merchantability and fitness for a particular purpose.

The product includes ACE(TM) and TAO(TM) software copyrighted by Douglas C. Schmidt and his research group at Washington University, University of California, Irvine, and Vanderbilt University, Copyright (©) 1993-2006, all rights reserved.

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (copyright The OpenSSL Project. All Rights Reserved) and redistribution of this software is subject to terms available at <http://www.openssl.org> and <http://www.openssl.org/source/license.html>.

This product includes Curl software which is Copyright 1996-2013, Daniel Stenberg, <daniel@haxx.se>. All Rights Reserved. Permissions and limitations regarding this software are subject to terms available at <http://curl.haxx.se/docs/copyright.html>. Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

The product includes software copyright 2001-2005 (©) MetaStuff, Ltd. All Rights Reserved. Permissions and limitations regarding this software are subject to terms available at <http://www.dom4j.org/license.html>.

The product includes software copyright © 2004-2007, The Dojo Foundation. All Rights Reserved. Permissions and limitations regarding this software are subject to terms available at <http://dojotoolkit.org/license>.

This product includes ICU software which is copyright International Business Machines Corporation and others. All rights reserved. Permissions and limitations regarding this software are subject to terms available at <http://source.icu-project.org/repos/icu/icu/trunk/license.html>.

This product includes software copyright © 1996-2006 Per Bothner. All rights reserved. Your right to use such materials is set forth in the license which may be found at <http://www.gnu.org/software/kawa/Software-License.html>.

This product includes OSSP UUID software which is Copyright © 2002 Ralf S. Engelschall, Copyright © 2002 The OSSP Project Copyright © 2002 Cable & Wireless Deutschland. Permissions and limitations regarding this software are subject to terms available at <http://www.opensource.org/licenses/mit-license.php>.

This product includes software developed by Boost (<http://www.boost.org/>) or under the Boost software license. Permissions and limitations regarding this software are subject to terms available at http://www.boost.org/LICENSE_1_0.txt.

This product includes software copyright © 1997-2007 University of Cambridge. Permissions and limitations regarding this software are subject to terms available at <http://www.pcre.org/license.txt>.

This product includes software copyright © 2007 The Eclipse Foundation. All Rights Reserved. Permissions and limitations regarding this software are subject to terms available at <http://www.eclipse.org/org/documents/epl-v10.php> and at <http://www.eclipse.org/org/documents/edl-v10.php>.

This product includes software licensed under the terms at <http://www.tcl.tk/software/tcltk/license.html>, <http://www.bosrup.com/web/overlib/?License>, <http://www.stlport.org/doc/license.html>, <http://asm.ow2.org/license.html>, <http://www.cryptix.org/LICENSE.TXT>, <http://hsqldb.org/web/hsqLicense.html>, <http://httpunit.sourceforge.net/doc/license.html>, <http://jung.sourceforge.net/license.txt>, http://www.gzip.org/zlib/zlib_license.html, <http://www.openldap.org/software/release/license.html>, <http://www.libssh2.org>, <http://slf4j.org/license.html>, <http://www.sente.ch/software/OpenSourceLicense.html>, <http://fusesource.com/downloads/license-agreements/fuse-message-broker-v-5-3-license-agreement>, <http://antlr.org/license.html>, <http://aopalliance.sourceforge.net/>, <http://www.bouncycastle.org/license.html>, <http://www.jgraph.com/jgraphdownload.html>, <http://www.jcraft.com/jsch/LICENSE.txt>, http://jotm.objectweb.org/bsd_license.html, <http://www.w3.org/Consortium/Legal/2002/copyright-software-20021231>, <http://www.slf4j.org/license.html>, <http://nanoxml.sourceforge.net/orig/copyright.html>, <http://www.json.org/license.html>, <http://forge.ow2.org/projects/javaservice/>, <http://www.postgresql.org/about/licence.html>, <http://www.sqlite.org/copyright.html>, <http://www.tcl.tk/software/tcltk/license.html>, <http://www.jaxen.org/faq.html>, <http://www.jdom.org/docs/faq.html>, <http://www.slf4j.org/license.html>, <http://www.iodbc.org/dataspace/iodbc/wiki/IODBC/License>, <http://www.keplerproject.org/md5/license.html>, <http://www.toedter.com/en/jcalendar/license.html>, <http://www.edankert.com/bounce/index.html>, <http://www.net-snmp.org/about/license.html>, <http://www.openmdx.org/#FAQ>, http://www.php.net/license/3_01.txt, <http://srp.stanford.edu/license.txt>;

<http://www.schneier.com/blowfish.html>; <http://www.jmock.org/license.html>; <http://xsom.java.net>; <http://benalman.com/about/license/>; <https://github.com/CreateJS/EaselJS/blob/master/src/easeljs/display/Bitmap.js>; <http://www.h2database.com/html/license.html#summary>; <http://jsoncpp.sourceforge.net/LICENSE>; <http://jdbc.postgresql.org/license.html>; <http://protobuf.googlecode.com/svn/trunk/src/google/protobuf/descriptor.proto>; <https://github.com/rantav/hector/blob/master/LICENSE>; <http://web.mit.edu/Kerberos/krb5-current/doc/mitK5license.html>; <http://jibx.sourceforge.net/jibx-license.html>; <https://github.com/lyokato/libgeohash/blob/master/LICENSE>; <https://github.com/hjiang/jsonxx/blob/master/LICENSE>; <https://code.google.com/p/lz4/>; <https://github.com/jedisct1/libsodium/blob/master/LICENSE>; <http://one-jar.sourceforge.net/index.php?page=documents&file=license>; <https://github.com/EsotericSoftware/kryo/blob/master/license.txt>; <http://www.scala-lang.org/license.html>; <https://github.com/tinkerpops/blueprints/blob/master/LICENSE.txt>; <http://gee.cs.oswego.edu/dl/classes/EDU/oswego/cs/dl/util/concurrent/intro.html>; <https://aws.amazon.com/asl/>; <https://github.com/twbs/bootstrap/blob/master/LICENSE>; <https://sourceforge.net/p/xmlunit/code/HEAD/tree/trunk/LICENSE.txt>; <https://github.com/documentcloud/underscore-contrib/blob/master/LICENSE>, and <https://github.com/apache/hbase/blob/master/LICENSE.txt>.

This product includes software licensed under the Academic Free License (<http://www.opensource.org/licenses/afl-3.0.php>), the Common Development and Distribution License (<http://www.opensource.org/licenses/cddl1.php>), the Common Public License (<http://www.opensource.org/licenses/cpl1.0.php>), the Sun Binary Code License Agreement Supplemental License Terms, the BSD License (<http://www.opensource.org/licenses/bsd-license.php>), the new BSD License (<http://opensource.org/licenses/BSD-3-Clause>), the MIT License (<http://www.opensource.org/licenses/mit-license.php>), the Artistic License (<http://www.opensource.org/licenses/artistic-license-1.0>) and the Initial Developer's Public License Version 1.0 (<http://www.firebirdsql.org/en/initial-developer-s-public-license-version-1-0/>).

This product includes software copyright © 2003-2006 Joe Walnes, 2006-2007 XStream Committers. All rights reserved. Permissions and limitations regarding this software are subject to terms available at <http://xstream.codehaus.org/license.html>. This product includes software developed by the Indiana University Extreme! Lab. For further information please visit <http://www.extreme.indiana.edu/>.

This product includes software Copyright (c) 2013 Frank Balluffi and Markus Moeller. All rights reserved. Permissions and limitations regarding this software are subject to terms of the MIT license.

See patents at <https://www.informatica.com/legal/patents.html>.

DISCLAIMER: Informatica LLC provides this documentation "as is" without warranty of any kind, either express or implied, including, but not limited to, the implied warranties of noninfringement, merchantability, or use for a particular purpose. Informatica LLC does not warrant that this software or documentation is error free. The information provided in this software or documentation may include technical inaccuracies or typographical errors. The information in this software and documentation is subject to change at any time without notice.

NOTICES

This Informatica product (the "Software") includes certain drivers (the "DataDirect Drivers") from DataDirect Technologies, an operating company of Progress Software Corporation ("DataDirect") which are subject to the following terms and conditions:

1. THE DATADIRECT DRIVERS ARE PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT.
2. IN NO EVENT WILL DATADIRECT OR ITS THIRD PARTY SUPPLIERS BE LIABLE TO THE END-USER CUSTOMER FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, CONSEQUENTIAL OR OTHER DAMAGES ARISING OUT OF THE USE OF THE ODBC DRIVERS, WHETHER OR NOT INFORMED OF THE POSSIBILITIES OF DAMAGES IN ADVANCE. THESE LIMITATIONS APPLY TO ALL CAUSES OF ACTION, INCLUDING, WITHOUT LIMITATION, BREACH OF CONTRACT, BREACH OF WARRANTY, NEGLIGENCE, STRICT LIABILITY, MISREPRESENTATION AND OTHER TORTS.

The information in this documentation is subject to change without notice. If you find any problems in this documentation, please report them to us in writing at Informatica LLC 2100 Seaport Blvd. Redwood City, CA 94063.

INFORMATICA LLC PROVIDES THE INFORMATION IN THIS DOCUMENT "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT.

Publication Date: 2018-09-28

Table of Contents

Preface	6
Informatica Resources.	6
Informatica Network.	6
Informatica Knowledge Base.	6
Informatica Documentation.	6
Informatica Product Availability Matrixes.	7
Informatica Velocity.	7
Informatica Marketplace.	7
Informatica Global Customer Support.	7
 Chapter 1: Introduction to PowerExchange for Vertica.....	 8
PowerExchange for Vertica Overview.	8
Introduction to Vertica.	8
 Chapter 2: PowerExchange for Vertica Configuration.....	 10
PowerExchange for Vertica Configuration Overview.	10
Prerequisites.	11
Vertica Configuration File for Linux.	11
Configuration Options.	11
Sample Vertica Configuration File.	12
Vertica Configuration for Windows.	12
PowerExchange for Vertica Plug-in Registration.	12
Upgrading PowerExchange for Vertica.	13
 Chapter 3: Vertica Sources and Targets.....	 14
Vertica Sources and Targets Overview.	14
Configuring ODBC Data Sources on Windows.	14
Configuring ODBC Data Sources on Linux.	15
Importing Vertica Source and Target Definitions.	15
 Chapter 4: Vertica Mappings.....	 17
Vertica Mappings Overview.	17
Vertica Mapping Example.	17
 Chapter 5: Vertica Sessions.....	 19
Vertica Sessions Overview.	19
PowerExchange for Vertica Connections.	20
Vertica Relational Connection Properties.	20
Vertica Bulk Connection Properties.	21
Configuring a Vertica Connection.	21

Copy Method.	22
Exception File.	22
Rejected File.	23
Non-Data Nodes.	24
Staging File and Compression.	24
Session Configuration for Vertica Targets in Bulk Mode.	24
Partitioning for Vertica Targets.	26
Appendix A: Vertica Data Type Reference.	28
Data Type Reference Overview.	28
Vertica and Transformation Data Types.	28
Index.	31

Preface

The *Informatica PowerExchange® for Vertica User Guide for PowerCenter®* describes how to read data from and write data to a Vertica database. The guide is written for database administrators and developers who are responsible for reading data from or writing data to a Vertica database. This guide assumes that you have knowledge of relational database concepts and database engines, Vertica, and PowerCenter.

Informatica Resources

Informatica Network

Informatica Network hosts Informatica Global Customer Support, the Informatica Knowledge Base, and other product resources. To access Informatica Network, visit <https://network.informatica.com>.

As a member, you can:

- Access all of your Informatica resources in one place.
- Search the Knowledge Base for product resources, including documentation, FAQs, and best practices.
- View product availability information.
- Review your support cases.
- Find your local Informatica User Group Network and collaborate with your peers.

Informatica Knowledge Base

Use the Informatica Knowledge Base to search Informatica Network for product resources such as documentation, how-to articles, best practices, and PAMs.

To access the Knowledge Base, visit <https://kb.informatica.com>. If you have questions, comments, or ideas about the Knowledge Base, contact the Informatica Knowledge Base team at KB_Feedback@informatica.com.

Informatica Documentation

To get the latest documentation for your product, browse the Informatica Knowledge Base at https://kb.informatica.com/_layouts/ProductDocumentation/Page/ProductDocumentSearch.aspx.

If you have questions, comments, or ideas about this documentation, contact the Informatica Documentation team through email at infa_documentation@informatica.com.

Informatica Product Availability Matrixes

Product Availability Matrixes (PAMs) indicate the versions of operating systems, databases, and other types of data sources and targets that a product release supports. If you are an Informatica Network member, you can access PAMs at

<https://network.informatica.com/community/informatica-network/product-availability-matrices>.

Informatica Velocity

Informatica Velocity is a collection of tips and best practices developed by Informatica Professional Services. Developed from the real-world experience of hundreds of data management projects, Informatica Velocity represents the collective knowledge of our consultants who have worked with organizations from around the world to plan, develop, deploy, and maintain successful data management solutions.

If you are an Informatica Network member, you can access Informatica Velocity resources at <http://velocity.informatica.com>.

If you have questions, comments, or ideas about Informatica Velocity, contact Informatica Professional Services at ips@informatica.com.

Informatica Marketplace

The Informatica Marketplace is a forum where you can find solutions that augment, extend, or enhance your Informatica implementations. By leveraging any of the hundreds of solutions from Informatica developers and partners, you can improve your productivity and speed up time to implementation on your projects. You can access Informatica Marketplace at <https://marketplace.informatica.com>.

Informatica Global Customer Support

You can contact a Global Support Center by telephone or through Online Support on Informatica Network.

To find your local Informatica Global Customer Support telephone number, visit the Informatica website at the following link:

<http://www.informatica.com/us/services-and-training/support-services/global-support-centers>.

If you are an Informatica Network member, you can use Online Support at <http://network.informatica.com>.

CHAPTER 1

Introduction to PowerExchange for Vertica

This chapter includes the following topics:

- [PowerExchange for Vertica Overview, 8](#)
- [Introduction to Vertica, 8](#)

PowerExchange for Vertica Overview

PowerExchange for Vertica provides connectivity between PowerCenter and Vertica.

You can use relational mode or bulk mode to process Vertica data.

In bulk mode, you can use PowerExchange for Vertica to write large volumes of data to a Vertica database. You can then run queries on the Vertica database to analyze the data. You use an ODBC data source to connect to the Vertica database. When you run a Vertica session in bulk mode, the PowerCenter Integration Service creates a pipe, runs the copy command, and writes the data to a pipe. The copy command reads the data from the pipe and writes the data to the Vertica target.

You can also use relational mode to read or write data. Use relational mode when you want to read large volumes of data from a Vertica source or when you want to update or delete data in a Vertica target.

Example

You work for a social media organization that wants to track user activity on the website. You can use a Vertica database to store the user online activity. You can then query the Vertica database and use it for analysis. For example, you can use PowerExchange for Vertica to find where your top users are located.

Introduction to Vertica

The Vertica database is a distributed storage system that you can use to store and analyze terabyte to petabytes of data.

Unlike traditional databases, Vertica uses a column-oriented approach for storing data. It stores each column separately on the disk. When you run a query on a Vertica database, only the columns that are called in the query will be accessed. The columns that are not called in the query will be ignored. Therefore, you can

process queries very efficiently since the amount of data accessed is much smaller than a traditional row store.

Vertica also uses encoding, sorting, and compression while storing data.

CHAPTER 2

PowerExchange for Vertica Configuration

This chapter includes the following topics:

- [PowerExchange for Vertica Configuration Overview, 10](#)
- [Prerequisites, 11](#)
- [Vertica Configuration File for Linux, 11](#)
- [Vertica Configuration for Windows, 12](#)
- [PowerExchange for Vertica Plug-in Registration, 12](#)
- [Upgrading PowerExchange for Vertica, 13](#)

PowerExchange for Vertica Configuration Overview

PowerExchange for Vertica is installed with PowerCenter.

You must configure PowerExchange for Vertica before you read data from or write data to a Vertica database. You can configure PowerExchange for Vertica on Windows and Linux operating systems.

To configure PowerExchange for Vertica, perform the following steps:

1. Complete the prerequisites.
2. Define the configuration options that the Vertica ODBC driver must use based on the operating system you use.
 - If you use Linux operating systems, create a configuration file named `vertica.ini` and define the Vertica configuration options in it. Define the `VERTICAINI` environment variable and set it to the path where you stored the `vertica.ini` configuration file.
 - If you use Windows operating systems, define the configuration options in the ODBC data source that you create to access the Vertica database.
3. If you use bulk mode to write data, register the PowerExchange for Vertica plug-in with the PowerCenter repository.

Prerequisites

Before you configure PowerExchange for Vertica, perform the following prerequisite tasks:

1. Install PowerCenter.
2. Install the Vertica database.
3. Install the Vertica ODBC driver to import Vertica source and target definitions. You can download the Vertica ODBC driver from the following site:
<https://my.vertica.com/download/vertica/client-drivers/>

Vertica Configuration File for Linux

If you use Linux operating systems, the Vertica ODBC driver uses a Vertica configuration file to determine the configuration options that it must use while connecting to a Vertica database.

Use a text editor to create the Vertica configuration file and name the file as `vertica.ini`. In the `vertica.ini` file, define the configuration options that you want the Vertica ODBC driver to use to connect to a Vertica database.

After you configure the `vertica.ini` file, define an environment variable called `VERTICAINI` and set it to the path where you stored the `vertica.ini` configuration file.

Configuration Options

The configuration options that you define in the `vertica.ini` file determine how the Vertica ODBC driver connects to the Vertica database.

DriverManagerEncoding

Defines the UTF encoding standard that the driver manager must use to connect to the Vertica database.

Set this option to UTF-16 for the Vertica ODBC driver.

ODBCInstLib

Defines the path where you stored the installer library file for the Vertica ODBC driver.

Set this option to the file path where you stored the `libodbcinst.so` file.

ErrorMessagesPath

Defines the path where the localized error message files for the Vertica ODBC driver are stored. These files are stored in the same directory as the Vertica ODBC driver files.

This option is required. If you do not set it and the ODBC driver encounters an error, the ODBC driver prints an error message stating that the `ODBCMessages.xml` file is missing.

LogLevel

Defines the minimum severity for the Vertica ODBC driver to log messages in the log file. These messages are specific to the connection errors that the Vertica ODBC driver encounters.

Set this option to one of the following values:

- 0. No logging

- 1. Fatal errors
- 2. Errors
- 3. Warnings
- 4. Informational messages
- 5. Debug messages
- 6. Trace, all messages

For example, if you set the value to 4, the Vertica ODBC client driver logs all informational messages, warnings, errors, and fatal errors.

Note: The LogLevel value does not affect the tracing level session property that you specify in the PowerCenter session properties. The tracing level session property determines the amount of detail that the PowerCenter Integration Service writes in the session log.

LogPath

Defines the path where the Vertica ODBC driver must create the log file.

Sample Vertica Configuration File

You can use the sample Vertica configuration file as a guideline to create the file.

```
[VerticaDriverName]
DriverManagerEncoding=UTF-16
ODBCInstLib=<Informatica install Dir>/ODBC7.1/lib/libodbcinst.so
ErrorMessagesPath=/opt/vertica/lib64
LogLevel=4
LogPath=/tmp
```

Vertica Configuration for Windows

If you use Windows operating systems, define the configuration options in the ODBC data source that you create to access the Vertica database.

For information about Vertica configuration for Windows, see [“Configuring ODBC Data Sources on Windows” on page 14](#).

PowerExchange for Vertica Plug-in Registration

To write data to a Vertica database in bulk mode, you must register the PowerExchange for Vertica plug-in with the PowerCenter repository. You do not need to register the plug-in with the repository to read or write data in relational mode. When you upgrade PowerCenter from an earlier version and want to write data in bulk mode, you must update the PowerExchange for Vertica plug-in registration.

A plug-in is an XML file that defines the functionality of PowerExchange for Vertica. The plug-in file for PowerExchange for Vertica is named `VerticaConnector.xml`. When you install PowerCenter, the installer copies the `VerticaConnector.xml` file to the following directory:

```
<Informatica Installation Directory>\server\bin\Plugin
```

To register the plug-in, the repository must be running in exclusive mode. Use the Administrator tool or the `pmrep RegisterPlugin` command to register the plug-in.

Note: If you do not have the Manage Services privilege to register the plug-in, contact the administrator who manages the PowerCenter Repository Service.

Upgrading PowerExchange for Vertica

You can upgrade PowerExchange for Vertica to a later version to take advantage of new features. After you upgrade PowerExchange for Vertica, you must update the existing plug-in file registration to write data in bulk mode.

1. Upgrade PowerCenter.
2. Start the PowerCenter Repository Service in exclusive mode.
To change the PowerCenter Repository Service operating mode, you can use the Administrator tool or the `infacmd UpdateRepositoryService` command.
3. Run the `pmrep connect` command to connect to the PowerCenter Repository Service.
4. Run the `pmrep upgradeMMDTablesForVertica` command to upgrade PowerExchange for Vertica.
Enter the following command:

```
pmrep upgradeMMDTablesForVertica
```
5. In Informatica Administrator, click the **Plug-Ins** tab.
6. Click the **Register Plug-in** icon.
7. Browse and select the Vertica plug-in file and select the option to update the existing plug-in file registration.
8. Enter your repository user name and password.
9. Click **OK**.
The PowerCenter Repository Service updates the existing Vertica plug-in file registration.
10. Restart the PowerCenter Repository Service to run in normal mode.

CHAPTER 3

Vertica Sources and Targets

This chapter includes the following topics:

- [Vertica Sources and Targets Overview, 14](#)
- [Configuring ODBC Data Sources on Windows, 14](#)
- [Configuring ODBC Data Sources on Linux, 15](#)
- [Importing Vertica Source and Target Definitions, 15](#)

Vertica Sources and Targets Overview

Vertica source definitions and target definitions represent metadata for Vertica tables.

You can import Vertica source and target definitions. You can also create Vertica target definitions based on a source definition, or manually create target definitions.

Before you import a Vertica source definition or target definition, you must create an ODBC data source for the Vertica database.

Configuring ODBC Data Sources on Windows

On Windows operating systems, use the ODBC Data Source Administrator to configure an ODBC data source for each Vertica database that you want to access. Use the Vertica ODBC driver to create a DSN connection to the Vertica database.

1. Start the ODBC Data Source Administrator.
 - On a Windows 32-bit system, open the **Control Panel** and click **Administrative Tools**. Then, double-click **Data Sources (ODBC)**.
 - On a Windows 64-bit system, enter the following command at the command prompt:
`%windir%\SysWOW64\odbcad32.exe`

The **ODBC Data Source Administrator** dialog box appears.

2. Click the **System DSN** tab, and then click **Add**.

The **Create New Data Source** dialog box appears.

3. Select the Vertica ODBC driver from the list of available drivers and click **Finish**.

An entry for the Vertica ODBC driver is added under the **System Data Sources** section.

4. On the **System DSN** tab, select the entry for the Vertica ODBC driver, and click **Configure**.
The **HP Vertica ODBC DSN Configuration** dialog box appears.
5. Enter a DSN name.
6. Enter the server name, port number, and the name of the Vertica database that you want to connect to.
7. Enter the user name and password to connect to the Vertica database.
8. Click **Test connection** to verify if the connection is successful.
9. Click **OK** to save the changes.
10. Configure additional options in the ODBC Data Source Administrator as needed.

Configuring ODBC Data Sources on Linux

On Linux operating systems, update the `odbc.ini` file entries to configure an ODBC data source for each Vertica database that you want to access.

1. Locate the `odbc.ini` file in the following directory: `<Informatica Installation Directory>\ODBC7.1\odbc.ini`
2. In the `odbc.ini` file, add entries for all the Vertica databases that you want to connect to.

The following section shows a sample Vertica entry in the `odbc.ini` file:

```
[ODBC Data Sources]
vertica_odbc=libverticaodbc.so
[ODBC]
InstallDir=/export/home/vertica/Informatica/9.6.1/ODBC7.1
Trace=0
TraceFile=/export/home/vertica/Informatica/9.6.1/odbctrace.out
TraceDll=/export/home/vertica/Informatica/9.6.1/ODBC7.1/lib/DWtrc27.so
[vertica_odbc]
Description = Vmart Database
Driver = /opt/vertica/lib64/libverticaodbc.so
Database = <database_name>
Servername = <server_name>
UID= <user_ID>
PWD= <password>
Port = <port_number>
ConnSettings=
SSLKeyFile=
SSLCertFile=
Locale=UTF-8
```

3. Set the ODBCINI environment variable to the path where the `odbc.ini` file is stored.

Importing Vertica Source and Target Definitions

To import a Vertica source or target definition, use an ODBC data source.

1. In the Source Analyzer, click **Sources > Import from Database**. Alternatively, in the Target Designer, click **Targets > Import from Database**.
The **Import Tables** dialog box appears.
2. Select the ODBC data source that you created for the Vertica database.

3. Enter the user name and password to connect to the Vertica database.
If you are not the owner of the table that you want to use, specify the owner name.
4. Click **Connect**.
A list of Vertica tables appears.
5. Drill down through the list of database objects to view the available tables.
6. Select the Vertica table that you want to import into the PowerCenter repository.
You can press and hold the Shift key to select a block of tables, or press and hold the Ctrl key to make non-contiguous selections. You can also use the **Select all** and **Select none** buttons to select or clear all available targets.
7. Click **OK**.
If you imported a Vertica source, the selected source definition appears under the Sources node in the Navigator. If you imported a Vertica target, the selected target definition appears under the Targets node in the Navigator.
The database type in the source and target definitions appears as Vertica.

CHAPTER 4

Vertica Mappings

This chapter includes the following topics:

- [Vertica Mappings Overview, 17](#)
- [Vertica Mapping Example, 17](#)

Vertica Mappings Overview

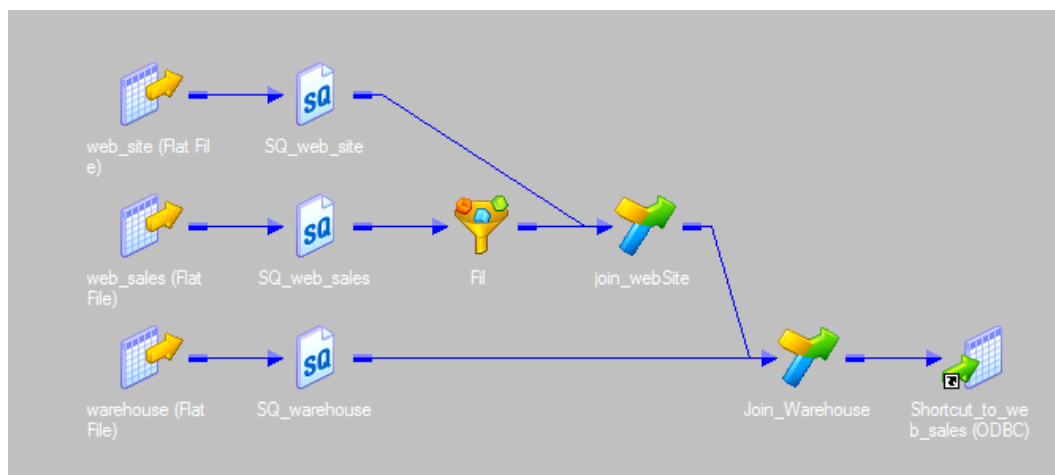
After you import a Vertica source or target definition into the PowerCenter repository, you can create a mapping to read data from a Vertica source or write data to a Vertica target.

You can read data from multiple Vertica sources and write data to multiple Vertica targets.

Vertica Mapping Example

Your organization sells electronic products through multiple websites and warehouses. You want to determine the sales details for a particular warehouse and write the sales information to a Vertica database.

The following image shows the Vertica mapping example:



The Vertica mapping contains the following objects:

Sources

The mapping contains the following flat file sources:

- `web_site`. Contains details of the websites through which customers can place orders for products.
- `web_sales`. Contains details of the sales done through websites and warehouses.
- `warehouse`. Contains details of the warehouses from where the products are shipped.

Transformations

The `Flt` filter transformation filters the sales details for a particular warehouse.

The `join_webSite` Joiner transformation joins the sales details with the website details. The sales record contains the website key. The transformation uses the website key to retrieve website details from the `web_site` source.

The `Join_Warehouse` Joiner transformation joins the output of the `join_webSite` Joiner transformation with the warehouse details.

Target

The target named `Web_sales` is a Vertica database. Specify the columns in the Vertica table to which you want to write the data.

When you run the mapping, the PowerCenter Integration Service writes the sales information of the specified warehouse to the Vertica target database. You can then run queries on the Vertica target database to analyze the data.

CHAPTER 5

Vertica Sessions

This chapter includes the following topics:

- [Vertica Sessions Overview, 19](#)
- [PowerExchange for Vertica Connections, 20](#)
- [Copy Method, 22](#)
- [Exception File, 22](#)
- [Rejected File, 23](#)
- [Non-Data Nodes, 24](#)
- [Staging File and Compression, 24](#)
- [Session Configuration for Vertica Targets in Bulk Mode, 24](#)
- [Partitioning for Vertica Targets, 26](#)

Vertica Sessions Overview

After you create a Vertica mapping in the Designer, you can create a session and configure the session to read data from or write data to a Vertica database. You can also run sessions on a grid.

You can use relational mode or bulk mode to process Vertica data. To read data, you must use relational mode. To write data, you can use relational mode or bulk mode.

Relational Mode

Use relational mode when you want to read large volumes of data from a Vertica source or when you want to update or delete data in a Vertica target.

To read data in relational mode, you must create a Vertica relational connection and configure the session to use a relational reader. Then, define the session properties for the relational reader. The session properties for relational mode are the same as that of relational sources.

Similarly, to write data in relational mode, you must create a Vertica relational connection and configure the session to use a relational writer. Then, define the session properties for the relational writer. The session properties for relational mode are the same as that of relational targets.

Bulk Mode

Use bulk mode when you want to insert large volumes of data into a Vertica target.

To write data in bulk mode, you must create a Vertica bulk connection and configure the session to use a Vertica bulk writer. Then, define the session properties that the copy command must use to write data to the Vertica target.

PowerExchange for Vertica Connections

You must create a Vertica connection for each Vertica source or target that you want to access. When you configure a Vertica connection, you specify the connection attributes that the PowerCenter Integration Service uses to connect to Vertica.

The type of Vertica connection that you need to create depends on the mode that you use to read or write data. To read or write data in relational mode, you must create a Vertica relational connection. To write data in bulk mode, you must create a Vertica bulk connection.

Vertica Relational Connection Properties

The following table describes the properties that you must configure for a Vertica relational connection:

Property	Description
Name	Enter a name for the connection.
Type	The connection type is set by default. You cannot edit this value.
User Name	Enter the user name to connect to the Vertica database. The user must have read permissions on the schema where the Vertica source table resides. Similarly, the user must have write permissions on the schema where the Vertica target table resides.
Password	Enter the password to connect to the Vertica database.
Use Parameter in Password	Indicates that the password for the repository user name is a session parameter, <code>\$ParamName</code> . Define the password in the workflow or session parameter file, and encrypt it by using the <code>pmpasswd CRYPT_DATA</code> option. Default is disabled.
Connect String	Enter the name of the ODBC data source that you created for the Vertica database.
Code Page	Select the code page that the PowerCenter Integration Service must use to read or write data.
Connection Environment SQL	Runs an SQL command with each database connection. Default is disabled.
Transaction Environment SQL	Runs an SQL command before the initiation of each transaction. Default is disabled.
Connection Retry Period	The number of seconds that the PowerCenter Integration Service attempts to reconnect to the database if the connection fails. If the PowerCenter Integration Service cannot connect to the database in the retry period, the session fails. Default value is 0.
Load Balance	Select this option to enable load balancing on the Vertica database. When you enable load balancing, the PowerCenter Integration Service requests a load balanced connection with the Vertica database. Default is enabled. For more information on load balancing, see the Vertica database documentation.

Vertica Bulk Connection Properties

The following table describes the properties that you must configure for a Vertica bulk connection:

Property	Description
Name	Enter a name for the connection.
Type	The connection type is set by default. You cannot edit this value.
User Name	Enter the user name to connect to the Vertica database. The user must have write permissions on the schema where the Vertica target table resides.
Password	Enter the password to connect to the Vertica database.
Use Parameter in Password	Indicates that the password for the repository user name is a session parameter, <code>\$ParamName</code> . Define the password in the workflow or session parameter file, and encrypt it by using the <code>mpasswd CRYPT_DATA</code> option. Default is disabled.
Connect String	Enter the name of the ODBC data source that you created for the Vertica target.
Code Page	The code page is set by default. You cannot edit this value.
Schema	Enter the name of the schema where the Vertica target table resides. Default is public.
Load Balance	Select this option to enable load balancing on the Vertica database. When you enable load balancing, the PowerCenter Integration Service requests a load balanced connection with the Vertica database. Default is enabled. For more information on load balancing, see the Vertica database documentation.

Configuring a Vertica Connection

Before you run a Vertica session, create a Vertica connection.

1. In the Workflow Manager, connect to a repository.
2. Click **Connections > Relational**.
The **Relational Connection Browser** dialog box appears.
3. Click **New**.
The **Select Subtype** dialog box appears.
4. Create a Vertica connection.
 - To create a Vertica relational connection, select **Vertica** from the **Select Subtype** list.
 - To create a Vertica bulk connection, select **Vertica Bulk Connection** from the **Select Subtype** list.
5. Click **OK**.
The **Connection Object Definition** dialog box appears.
6. Enter the connection properties.
7. Click **OK**.
The Vertica connection appears in the **Connection Browser** list.

Copy Method

The copy method defines how the copy command writes data to the Vertica target. It defines the type of container to which the copy command writes the data.

You can write data to the following types of containers in Vertica:

- Write Optimized Store in memory (WOS) container. WOS is a memory-resident data structure that enables high-speed data loading. WOS stores data without any compression or indexing.
- Read Optimized Store on disk (ROS) container. ROS is a highly optimized, read-oriented, disk storage structure that compresses and indexes data.

You can set the copy method to one of the following values:

Auto

The copy command writes the data continuously to a WOS container. When the WOS container is full, the copy command writes the data to a ROS container.

Select this method when you want to write multiple data sets that are lesser than 100 MB in size. Since the WOS container does not compress and index data, you can write multiple data sets at high speed and then consolidate the data sets into a ROS container.

Direct

The copy command directly writes the data to a ROS container without writing it first to a WOS container.

Select this method when you want to write data sets that are greater than 100 MB in size, at a time. Since the ROS container compresses and indexes data, you can reduce the disk storage and improve query execution.

Trickle

The copy command writes the data incrementally to a WOS container. When the WOS container is full, the copy command displays an error and rolls back the entire data load.

Select this method when you have written your initial data and want to write additional data on an ongoing basis. For example, you can write 10,000 to 100,000 rows per transaction into the WOS container.

Exception File

When you run a session in bulk mode, you can configure the PowerCenter Integration Service to log the copy command exceptions in an exception file. Use the exception file to understand why the exceptions occurred and take corrective action.

The exception file contains the following details:

- The record number that was rejected, along with the reason for the exception
- The file path and name of the rejected file that contains the rejected record
- The pipe location from where the copy command read the record

If you use partitions, the PowerCenter Integration Service creates an exception file for each partition.

By default, the PowerCenter Integration Service creates the exception file in the following path:

```
$PMRootDir/Temp/
```

If you want to change the default path, you can define the path in the session properties.

The PowerCenter Integration Service uses the following naming convention for the exception file:

```
Exceptions_PartNum_[Partition_Number]_[Table Name]_[Schema  
Name]_[Counter]_<milliseconds*1000>.dat
```

Sample Exception File

COPY: Input record 8 has been rejected (int8 out of range '-9223372036854775808' for column 1 (COL1_INTEGER)). See /export/qa_adp/PMATT_96/work_dir/Vertica/Linux.64_RH6.1/exec_dir/Temp/RejectedData_PartNum_0_V0001_INTEGER_TGT_adpqa_schema_0_805000.dat, record 1 for the rejected record. This record was read from V0001_INTEGER_TGT_adpqa_schema_0_pipe_0

COPY: Input record 9 has been rejected (int8 out of range '-9223372036854775808' for column 1 (COL1_INTEGER)). See /export/qa_adp/PMATT_96/work_dir/Vertica/Linux.64_RH6.1/exec_dir/Temp/RejectedData_PartNum_0_V0001_INTEGER_TGT_adpqa_schema_0_805000.dat, record 2 for the rejected record. This record was read from V0001_INTEGER_TGT_adpqa_schema_0_pipe_0

COPY: Input record 10 has been rejected (int8 out of range '-9223372036854775808' for column 1 (COL1_INTEGER)). See /export/qa_adp/PMATT_96/work_dir/Vertica/Linux.64_RH6.1/exec_dir/Temp/RejectedData_PartNum_0_V0001_INTEGER_TGT_adpqa_schema_0_805000.dat, record 3 for the rejected record. This record was read from V0001_INTEGER_TGT_adpqa_schema_0_pipe_0

Rejected File

When you run a session in bulk mode, you can configure the PowerCenter Integration Service to write the rejected records to a rejected file. The rejected file contains all the records that were rejected by the copy command and not written to the Vertica target.

The rejected file displays the record details of each column separated by a | character. Therefore, if you have sensitive data, you must encrypt the data before you run the session.

If you use partitions, the PowerCenter Integration Service creates a rejected file for each partition.

By default, the PowerCenter Integration Service creates the rejected file in the following path:

```
$PMRootDir/Temp/
```

If you want to change the default path, you can define the path in the session properties.

The PowerCenter Integration Service uses the following naming convention for the rejected file:

```
RejectedData_PartNum_[Partition_Number]_[Table Name]_[Schema  
Name]_[Counter]_<milliseconds*1000>.dat
```

Sample Rejected File

```
"-9223372036854775808"|" -9223372036854775808"|" -9223372036854775808"|" -9223372036854775808"|" -  
-9223372036854775808"|" -9223372036854775808"|" -9223372036854775808"|" -9223372036854775808"|" -  
9223372036854775808"|" -9223372036854775808"|" -9223372036854775808"|" -9223372036854775808"|" -9  
223372036854775808"|" -9223372036854775808"|" -9223372036854775808"|" -9223372036854775808"|" -92  
23372036854775808"|" -9223372036854775808"
```

Non-Data Nodes

When you run a session in bulk mode, you can configure the copy command to connect to a non-data node in the Vertica cluster. A non-data node is a high-performance node that is used for writing data.

Staging File and Compression

When you write large volumes of data to a Vertica target in bulk mode, you can configure the session to create a staging file. You can also choose to compress the data to improve the session performance.

If you configure the session to create a staging file, the PowerCenter Integration Service writes the data to the staging file and then runs the copy command. The copy command reads the data from the staging file and writes the data to the Vertica target.

On UNIX operating systems, when you enable file staging, you can also choose to compress the data in a GZIP format. By compressing the data, you can reduce the size of data that is transferred over the network and improve session performance.

You can specify the compression factor to determine the level to which the data must be compressed. You can enter a value in the range of 1 to 9.

A value of 1 indicates that the PowerCenter Integration Service takes minimal time to compress the data, but the data compression level is low. A value of 9 indicates that the PowerCenter Integration Service takes maximal time to compress the data, but the data compression level is high in this case. For more information about configuring the data compression factor, see the GNU operating system documentation.

Session Configuration for Vertica Targets in Bulk Mode

To write data in bulk mode, configure the session to use a Vertica bulk connection and Vertica bulk writer. The copy command uses the session properties to write data to the Vertica target.

You can configure the session properties for Vertica targets in the **Transformations** view on the **Mapping** tab. Define the properties for each target instance in the session.

Method

The method that the copy command uses to write data to the Vertica target.

The method is set to INSERT because you can only insert data to a Vertica target. You cannot update or delete data.

Copy Method

The copy method that the copy command uses to write data to the Vertica target.

You can specify the copy method as auto, direct, or trickle.

Default is auto.

Delimiter

The character that separates successive input fields.

Default is |.

Escape As

The character that treats special characters in the data as regular characters.

Default is \.

Null AS

The character that the PowerCenter Integration Service replaces null values with when it writes data to the pipe. When the copy command reads the data from the pipe and writes it to the Vertica target, it converts the Null AS value to NULL.

Default is \n.

Quote

The character within which each column value gets enclosed. This option helps improve the readability of the data.

Default is ".

Exception File Path

The file path where the PowerCenter Integration Service creates the exception file.

Default is \$PMRootDir/Temp/.

Rejected File Path

The file path where the PowerCenter Integration Service creates the rejected file.

Default is \$PMRootDir/Temp/.

Abort On Errors

Select this option to configure the PowerCenter Integration Service to fail the session when it encounters an error.

Truncate target option

Select this option to truncate the Vertica target table before the PowerCenter Integration Service writes the data to the target.

Windows pipe buffer size

The size that the PowerCenter Integration Service allocates to buffer data before it writes the data to the Vertica target.

Enter a value between 1 KB and 2096128 KB.

Default is 4 KB.

You might need to test different settings for optimal performance. This attribute is applicable if the Informatica servers run on Windows operating systems.

Pipe Location

The file path where the PowerCenter Integration Service creates the pipes that it uses for data transfer.

Default is \$PMRootDir/Temp/.

This attribute is applicable if the Informatica servers run on UNIX operating systems.

Reject Max Records

Enter the maximum number of records that the PowerCenter Integration Service can reject after which it must fail the session.

Connect Non-Data Nodes

Select this option to connect to a non-data node to improve performance. Non-data nodes are high-performance nodes that are used for data loading.

Enable File Staging

Select this option to create a staging file in a local directory.

Default is disabled.

Stage File Directory

The file path where the PowerCenter Integration Service creates the staging file.

Data Compression mode

On a UNIX operating system, use this option to compress the data before you write it to the staging file. The PowerCenter Integration Service uses this attribute only if you enable file staging. Otherwise, it ignores this attribute.

You can specify the data compression mode as GZIP.

GZIP Compression factor

If you select the GZIP data compression mode, set a compression factor in the range of 1 to 9.

Default is 6.

Miscellaneous Attribute

Use this field to configure any Vertica attribute that is currently not available in the session properties.

Enter any attribute that Vertica supports along with the attribute value.

Use the following format:

```
ATTRIBUTE_NAME_1,ATTRIBUTE_VALUE_1;  
ATTRIBUTE_NAME_2,ATTRIBUTE_VALUE_2;  
<and so on>
```

Table Name Prefix

Enter a target owner name or prefix to override the table name prefix in the mapping.

Target Table Name

Enter a target table name to override the table name in the mapping.

Partitioning for Vertica Targets

If you need to write large amounts of data to a Vertica target in bulk mode, you can configure partitioning to improve session performance.

You can configure pass-through partitioning for Vertica sessions.

When you configure partitioning, the PowerCenter Integration Service runs multiple copy commands based on the number of partitions that you configured. For example, if you configured four partitions, the

PowerCenter Integration Service creates four pipes, and runs four copy commands. It then writes data into four pipes. Each copy command reads from one pipe and writes the data to the Vertica target.

If you configure partitioning, the PowerCenter Integration Service creates the exception file and the rejected file for each partition. Similarly, in the session log, the PowerCenter Integration Service prints the load statistics such as the applied rows and affected rows for each partition.

APPENDIX A

Vertica Data Type Reference

This appendix includes the following topics:

- [Data Type Reference Overview, 28](#)
- [Vertica and Transformation Data Types, 28](#)

Data Type Reference Overview

PowerCenter uses the following data types in Vertica mappings:

- Vertica native data types. Vertica data types appear in Vertica definitions in a mapping.
- Transformation data types. Set of data types that appear in the transformations. They are internal data types based on ANSI SQL-92 generic data types, which the PowerCenter Integration Service uses to move data across platforms. They appear in all transformations in a mapping.

When the PowerCenter Integration Service reads source data, it converts the native data types to the comparable transformation data types before transforming the data. When the PowerCenter Integration Service writes to target, it converts the transformation data types to the comparable native data types.

Vertica and Transformation Data Types

The following table compares the Vertica data types and the transformation data types:

Vertica Data Type	Transformation Data Type	Description
Bigint	Bigint	-9,223,372,036,854,775,808 to 9,223,372,036,854,775,807 Precision 19, scale 0
Binary	Binary	1 to 104,857,600 bytes
Boolean	String	1 to 104,857,600 characters
Bytea	Binary	1 to 104,857,600 bytes
Char	String	1 to 104,857,600 characters

Vertica Data Type	Transformation Data Type	Description
Date	Date/Time	Jan 1, 0001 A.D. to Dec 31, 9999 A.D. (precision to the nanosecond)
Datetime	Date/Time	Jan 1, 0001 A.D. to Dec 31, 9999 A.D. (precision to the nanosecond)
Decimal	Decimal	Precision 1 to 28, scale 0 to 28
Double precision	Double	Precision 15
Float	Double	Precision 15
Float (N)	Double	Precision 15
Float8	Double	Precision 15
Int	Bigint	-9,223,372,036,854,775,808 to 9,223,372,036,854,775,807 Precision 19, scale 0
Int8	Bigint	-9,223,372,036,854,775,808 to 9,223,372,036,854,775,807 Precision 19, scale 0
Integer	Bigint	-9,223,372,036,854,775,808 to 9,223,372,036,854,775,807 Precision 19, scale 0
Interval	String	1 to 104,857,600 characters
Money	Numeric (P, S)	Precision 1 to 28, scale 0 to 28
Number	Decimal	Precision 1 to 28, scale 0 to 28
Numeric	Decimal	Precision 1 to 28, scale 0 to 28
Raw	Binary	1 to 104,857,600 bytes
Real	Double	Precision 15
Smalldatetime	Date/Time	Jan 1, 0001 A.D. to Dec 31, 9999 A.D. (precision to the nanosecond)
Smallint	Bigint	-9,223,372,036,854,775,808 to 9,223,372,036,854,775,807 Precision 19, scale 0
Time	Date/Time	Jan 1, 0001 A.D. to Dec 31, 9999 A.D. (precision to the nanosecond)
Timestamp	Date/Time	Jan 1, 0001 A.D. to Dec 31, 9999 A.D. (precision to the nanosecond)
Tinyint	Bigint	-9,223,372,036,854,775,808 to 9,223,372,036,854,775,807 Precision 19, scale 0

Vertica Data Type	Transformation Data Type	Description
Varbinary	Binary	1 to 104,857,600 bytes
Varchar	String	1 to 104,857,600 characters

INDEX

C

configuring ODBC data sources
on Linux [15](#)
on Windows [14](#)

P

PowerExchange for Vertica
adapter overview [8](#)
configuration overview [10](#)
sessions overview [19](#)
upgrading [13](#)
PowerExchange for Vertica configuration
prerequisites [11](#)
PowerExchange for Vertica connections
bulk properties [21](#)
configuring [21](#)
relational properties [20](#)

R

registering plug-ins
PowerExchange for Vertica [12](#)

V

Vertica
configuration file [11](#)
configuration options [11](#)
data types overview [28](#)
mappings overview [17](#)
sample configuration file [12](#)
sources overview [14](#)
targets overview [14](#)
Vertica data types
mapping to transformation data types [28](#)
Vertica mappings
example [17](#)
Vertica session properties
configuration [24](#)
copy method [22](#)
exception file [22](#)
rejected file [23](#)
Vertica sources
importing [15](#)
Vertica targets
importing [15](#)
partitioning [26](#)