



Informatica® PowerExchange for Teradata
Parallel Transporter API
10.0

User Guide

© Copyright Informatica LLC 2012, 2018

This software and documentation contain proprietary information of Informatica LLC and are provided under a license agreement containing restrictions on use and disclosure and are also protected by copyright law. Reverse engineering of the software is prohibited. 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. This Software may be protected by U.S. and/or international Patents and other Patents Pending.

Use, duplication, or disclosure of the Software by the U.S. Government is subject to the restrictions set forth in the applicable software license agreement and as provided in DFARS 227.7202-1(a) and 227.7702-3(a) (1995), DFARS 252.227-7013(1)(ii) (OCT 1988), FAR 12.212(a) (1995), FAR 52.227-19, or FAR 52.227-14 (ALT III), as applicable.

The information in this product or documentation is subject to change without notice. If you find any problems in this product or documentation, please report them to us in writing.

Informatica, Informatica Platform, Informatica Data Services, PowerCenter, PowerCenterRT, PowerCenter Connect, PowerCenter Data Analyzer, PowerExchange, PowerMart, Metadata Manager, Informatica Data Quality, Informatica Data Explorer, Informatica B2B Data Transformation, Informatica B2B Data Exchange Informatica On Demand, Informatica Identity Resolution, Informatica Application Information Lifecycle Management, Informatica Complex Event Processing, Ultra Messaging and Informatica Master Data Management are trademarks or registered trademarks of Informatica LLC in the United States and in jurisdictions throughout the world. All 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 (c) 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/hsqldbLicense.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/licence.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.schneider.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/tinkerpop/blueprints/blob/master/LICENSE.txt>, <http://gee.cs.oswego.edu/dl/classes/EDU/oswego/cs/dl/util/concurrent/intro.html>, <https://aws.amazon.com/ssl/>, <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.

Publication Date: 2018-10-17

Table of Contents

Preface	6
Informatica Resources.	6
Informatica My Support Portal.	6
Informatica Documentation.	6
Informatica Product Availability Matrixes.	7
Informatica Web Site.	7
Informatica How-To Library.	7
Informatica Knowledge Base.	7
Informatica Support YouTube Channel.	7
Informatica Marketplace.	7
Informatica Velocity.	7
Informatica Global Customer Support.	8
 Chapter 1: Introduction to PowerExchange for Teradata Parallel Transporter API.....	 9
PowerExchange for Teradata Parallel Transporter API Overview.	9
 Chapter 2: PowerExchange for Teradata Parallel Transporter API Configuration.....	 10
Prerequisites.	10
Environment Variables.	11
Database Privileges.	12
JDBC Configuration.	12
TDCH Configuration for Teradata Mappings in a Hive Run-time Environment.	13
Configuration Parameters for Teradata Connector for Hadoop.	13
Step 1. Installing and Configuring Teradata Connector for Hadoop.	14
Step 2. Configuring the InfaTDCHConfig.txt File.	14
Step 3. Configuring the EnableTdch Custom Property.	15
 Chapter 3: PowerExchange for Teradata Parallel Transporter Connections... 	 16
Teradata Parallel Transporter Connection Overview.	16
Teradata Parallel Transporter Connection Properties.	16
Creating a Teradata Parallel Transporter Connection.	18
 Chapter 4: PowerExchange for Teradata Parallel Transporter API Data Objects.....	 20
Teradata Data Object Overview.	20
Teradata Data Object Sections.	21
Teradata Data Object Overview Properties.	21
Teradata Data Object Read Operation Properties.	22
Input Properties of the Data Object Read Operation.	22

Source Properties of the Data Object Read Operation.	24
Teradata Data Object Write Operation Properties.	24
Input Properties of the Data Object Write Operation.	25
Target Properties of the Data Object Write Operation.	29
Importing a Teradata Data Object.	30
Creating a Teradata Data Object Read Operation.	31
Creating a Teradata Data Object Write Operation.	31
 Chapter 5: PowerExchange for Teradata Parallel Transporter API Mappings..	33
Teradata Mappings Overview.	33
Hive Validation and Run-time Environments.	33
Teradata Mapping Example.	34
 Chapter 6: Teradata Parallel Transporter API Run-time Processing.....	35
Teradata Run-time Processing Overview.	35
System Operators.	35
Stream System Operator.	36
Load System Operator.	36
Update System Operator.	36
Error and Log Tables.	36
Error Tables.	37
Log Tables.	37
Row-level Error Processing.	38
Parameterization.	39
Partitioning.	40
Dynamic Partitioning.	41
 Appendix A: Data Type Reference.....	42
Data Type Reference Overview.	42
Teradata Data Types and Transformation Data Types.	42
 Index.	44

Preface

The *Informatica PowerExchange for Teradata Parallel Transporter API User Guide* provides information about reading data from and loading data to Teradata tables. It is written for database administrators and developers who create mappings to read data from or load data to Teradata.

This book assumes that you have knowledge of Teradata databases, Teradata Parallel Transporter, Informatica Developer, Informatica Big Data Management, and the database engines and systems in your environment.

Informatica Resources

Informatica My Support Portal

As an Informatica customer, the first step in reaching out to Informatica is through the Informatica My Support Portal at <https://mysupport.informatica.com>. The My Support Portal is the largest online data integration collaboration platform with over 100,000 Informatica customers and partners worldwide.

As a member, you can:

- Access all of your Informatica resources in one place.
- Review your support cases.
- Search the Knowledge Base, find product documentation, access how-to documents, and watch support videos.
- Find your local Informatica User Group Network and collaborate with your peers.

As a member, you can:

- Access all of your Informatica resources in one place.
- Search the Knowledge Base, find product documentation, access how-to documents, and watch support videos.
- Find your local Informatica User Group Network and collaborate with your peers.

Informatica Documentation

The Informatica Documentation team makes every effort to create accurate, usable documentation. If you have questions, comments, or ideas about this documentation, contact the Informatica Documentation team through email at infa_documentation@informatica.com. We will use your feedback to improve our documentation. Let us know if we can contact you regarding your comments.

The Documentation team updates documentation as needed. To get the latest documentation for your product, navigate to Product Documentation from <https://mysupport.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. You can access the PAMs on the Informatica My Support Portal at <https://mysupport.informatica.com>.

Informatica Web Site

You can access the Informatica corporate web site at <https://www.informatica.com>. The site contains information about Informatica, its background, upcoming events, and sales offices. You will also find product and partner information. The services area of the site includes important information about technical support, training and education, and implementation services.

Informatica How-To Library

As an Informatica customer, you can access the Informatica How-To Library at <https://mysupport.informatica.com>. The How-To Library is a collection of resources to help you learn more about Informatica products and features. It includes articles and interactive demonstrations that provide solutions to common problems, compare features and behaviors, and guide you through performing specific real-world tasks.

Informatica Knowledge Base

As an Informatica customer, you can access the Informatica Knowledge Base at <https://mysupport.informatica.com>. Use the Knowledge Base to search for documented solutions to known technical issues about Informatica products. You can also find answers to frequently asked questions, technical white papers, and technical tips. If you have questions, comments, or ideas about the Knowledge Base, contact the Informatica Knowledge Base team through email at KB_Feedback@informatica.com.

Informatica Support YouTube Channel

You can access the Informatica Support YouTube channel at <http://www.youtube.com/user/INFASupport>. The Informatica Support YouTube channel includes videos about solutions that guide you through performing specific tasks. If you have questions, comments, or ideas about the Informatica Support YouTube channel, contact the Support YouTube team through email at supportvideos@informatica.com or send a tweet to @INFASupport.

Informatica Marketplace

The Informatica Marketplace is a forum where developers and partners can share solutions that augment, extend, or enhance data integration implementations. By leveraging any of the hundreds of solutions available on the Marketplace, you can improve your productivity and speed up time to implementation on your projects. You can access Informatica Marketplace at <http://www.informaticamarketplace.com>.

Informatica Velocity

You can access Informatica Velocity at <https://mysupport.informatica.com>. 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 have questions, comments, or ideas about Informatica Velocity, contact Informatica Professional Services at ips@informatica.com.

Informatica Global Customer Support

You can contact a Customer Support Center by telephone or through the Online Support.

Online Support requires a user name and password. You can request a user name and password at <http://mysupport.informatica.com>.

The telephone numbers for Informatica Global Customer Support are available from the Informatica web site at <http://www.informatica.com/us/services-and-training/support-services/global-support-centers/>.

CHAPTER 1

Introduction to PowerExchange for Teradata Parallel Transporter API

This chapter includes the following topic:

- [PowerExchange for Teradata Parallel Transporter API Overview, 9](#)

PowerExchange for Teradata Parallel Transporter API Overview

PowerExchange for Teradata Parallel Transporter API integrates Informatica and Teradata Parallel Transporter API (Teradata PT API) to read and bulk load data. You can use PowerExchange for Teradata Parallel Transporter API to read data from and write large volumes of data into Teradata tables.

The Data Integration Service uses the Teradata PT API infrastructure to connect to Teradata. The Data Integration Service reads data using the Export system operator. The Export operator exports large amounts of data from Teradata tables or views. The Data Integration Service loads data in bulk using one of the following Teradata PT API system operators:

Load

Loads data in bulk into an empty Teradata table.

Stream

Performs insert, update, upsert, and delete operations against Teradata tables in near real-time.

Update

Performs insert, update, upsert, and delete operations against Teradata tables.

You can run Teradata mappings in native or Hive environments.

For example, in a native environment, you can store payroll information of XYZ organization for the past five years. You can use PowerExchange for Teradata Parallel Transporter API in native environment to write payroll information and then write to Teradata tables for storage.

For example, in a Hive environment, you can gather, store, and analyze large volumes of unstructured data like web logs. You can process these large amounts of data in a Hive environment and use PowerExchange for Teradata Parallel Transporter API to write meaningful data to Teradata tables for analysis.

CHAPTER 2

PowerExchange for Teradata Parallel Transporter API Configuration

This chapter includes the following topics:

- [Prerequisites, 10](#)
- [Environment Variables, 11](#)
- [Database Privileges, 12](#)
- [JDBC Configuration, 12](#)
- [TDCH Configuration for Teradata Mappings in a Hive Run-time Environment, 13](#)

Prerequisites

PowerExchange for Teradata Parallel Transporter API installs with the Informatica services and clients.

Before you use PowerExchange for Teradata Parallel Transporter API, you must perform the following tasks:

1. Install the Informatica services.
2. Install the Informatica clients. When you install the Informatica clients, the Developer tool is installed.
3. Create a Data Integration Service and a Model Repository Service in the Informatica domain.
4. Install Teradata Parallel Transporter.
 - To run Teradata mappings in a native environment, install Teradata Parallel Transporter on the machine that runs the Data Integration Service. You must update the Teradata Parallel Transporter installation to the latest patch.
 - To run Teradata mappings in a Hive run-time environment, install Teradata Parallel Transporter on every node of the Hadoop cluster. You must update the Teradata Parallel Transporter installation to the latest patch.
5. Copy the Teradata JDBC jars to the node that runs the Data Integration Service and update the CLASSPATH environment variable to include the Teradata JDBC jars on the client machine. Download the latest version of the Teradata JDBC driver from the following URL:
<http://downloads.teradata.com/download/connectivity/jdbc-driver>
6. To run Teradata mappings in a Hive run-time environment, install Informatica Big Data Management on every node of the Hadoop cluster.

7. Configure environment variables.

Environment Variables

You must configure Java and Teradata environment variables before you can use PowerExchange for Teradata Parallel Transporter API in a native or Hive environment.

In a native environment, set the environment variables on a node that runs the Data Integration Service. In a Hive environment, set the environment variables in the `hadoopEnv.properties` file. For more information, see the *Informatica Big Data Management User Guide*.

The following table describes the environment variables that you must configure on Windows and UNIX:

Environment Variable	Value
JAVA_HOME	<Informatica installation directory >/java
JRE_HOME	\${JAVA_HOME}/jre
<Shared Library Variable>	Set the shared library variable to include \$TWB_ROOT/lib64 or \$TWB_ROOT/lib for 32-bit files and \$TD_ICU_DATA/lib64 or \$TD_ICU_DATA/lib for 32-bit files.

Set the shared library environment variable based on the operating system. The following table describes the shared library variables for each operating system:

Operating System	Value
Windows	PATH
Linux	LD_LIBRARY_PATH

The following table describes the environment variables that you must configure on UNIX:

Environment Variable	Value
COPERR	Set to the location of the <code>errmsg.cat</code> file. For example: <code>/usr/lib</code>
COPLIB	Set to the location of the <code>clispb.dat</code> file. For example: <code>/usr/lib</code> .
TD_ICU_DATA	Set to the location of the Teradata ICU libraries. For example: <code>/opt/tdicu/</code>
THREADONOFF	On UNIX and Linux operating systems, set the THREADONOFF environment variable to 1 to enable multithreading support for Teradata Parallel Transporter processes.

Environment Variable	Value
TWB_ROOT	Set to the Teradata Parallel Transporter installation directory. For example: /opt/teradata/client/14.00/tbuild/ The 32-bit files are available in the following directory: \$TWB_ROOT/msg The 64-bit files are available in the following directory: \$TWB_ROOT/msg64 < TWB_ROOT>/lib or <TWB_ROOT>/lib64 should be added to the shared library value.
NLSPATH	Set to the following location: /opt/teradata/client/14.10/odbc_32/msg/%N:/opt/teradata/client/13.10/tbuild/msg64/%N:/usr/lib/nls/msg/%L/%N

Database Privileges

Verify read and write permissions on the Teradata database and select privileges on specific Data Dictionary tables, before you use PowerExchange for Teradata Parallel Transporter API.

Verify that you have the following database privileges:

- Read permission to read data from Teradata.
- Write permission to write data into a Teradata.
- Select privileges on DBC.Tables, DBC.Columns, DBC.UDTInfo, and DBC.Databases.
For information about select privileges, see the Teradata JDBC Driver documentation.

JDBC Configuration

Configure JDBC on the node that runs the Data Integration Service and on the client machine before you use PowerExchange for Teradata Parallel Transporter API.

Before you test a Teradata PT connection, perform the following tasks:

- To test a Teradata PT connection from the Administrator tool, copy Teradata JDBC jars to the node that runs the Data Integration Service, in the following location:

```
<Informatica Installation Directory>\<version folder>\services\shared\jars\thirdparty
```
- To test a Teradata PT connection from the client machine, update the CLASSPATH environment variable to include the Teradata JDBC jars on a client machine.

CLASSPATH is set to the location where the JDBC driver is installed:

```
<JDBC Driver Installation Directory>\Tera-JDBC\tdgssconfig.jar and <JDBC Driver Installation Directory>\Tera-JDBC\terajdbc4.jar
```

TDCH Configuration for Teradata Mappings in a Hive Run-time Environment

When you load data to a Teradata database in a Hive run-time environment, you can use the Teradata Connector for Hadoop (TDCH) Command Line Edition to increase performance. TDCH is a set of API and tools that Teradata Corporation provides for parallel processing of data between Teradata databases and the Hadoop ecosystem of products. You can download TDCH from the Teradata Developer Exchange website.

When you run a mapping in a Hive run-time environment, the Data Integration Service pushes the mapping to a Hadoop cluster and processes the mapping with one mapper task. When you enable TDCH to run a Teradata mapping in a Hive run-time environment, TDCH loads the data to the Teradata database by using multiple mapper tasks, which significantly increases the performance.

TDCH uses a configuration file named `InfaTDCHConfig.txt` to load the data. You can define configuration options in the `InfaTDCHConfig.txt` file to specify how you want TDCH to load the data.

To enable TDCH to run a Teradata mapping in a Hive run-time environment, perform the following tasks:

1. Install and configure TDCH.
2. Configure the `InfaTDCHConfig.txt` file.
3. Configure the `EnableTdch` custom property.

Configuration Parameters for Teradata Connector for Hadoop

Configuration parameters determine how TDCH loads data to Teradata.

Define the following configuration parameters in the `InfaTDCHConfig.txt` file:

libjars

Set the `libjars` parameter to the paths that contain the TDCH JAR files.

method

Set the `method` parameter to define the method that TDCH must use to load the data to the Teradata database.

You can set the `method` as `batch.insert` or `internal.fastload`.

nummappers

Set the `nummappers` parameter to define the maximum number of mapper tasks that TDCH must use while loading data to the Teradata database.

You can enter an integer value that is greater than or equal to zero. If you enter zero, the number of mapper tasks will be the same as the number of file blocks in HDFS.

Use either this parameter or the `numreducers` parameter, but not both.

For more information about the configuration options, see the Teradata Connector for Hadoop documentation.

Step 1. Installing and Configuring Teradata Connector for Hadoop

Install and configure TDCH to load data to Teradata in a Hive run-time environment.

Before you install and configure TDCH, you must create a user account in the Teradata Developer Exchange website so that you can download the TDCH installation package from the website.

1. Go to the Teradata Developer Exchange website: <http://developer.teradata.com/>
2. Download the Teradata Connector for Hadoop Command Line Edition installation package that is relevant for the Hadoop distribution that you use. For example, `teradata-connector-1.3.4-mapr3.1.noarch.rpm`.
3. Place the `teradata-connector.rpm` on all the nodes of the Hadoop cluster, and install the TDCH installation package.

The rpm package installs TDCH into the following directory: `/usr/lib/tdch/<version number folder>`

4. Access the following directory on the Informatica server: `<Informatica installation directory>\services\shared\hadoop\<Hadoop Distribution Name>\infaConf\`
5. Find the file named `hadoopEnv.properties`.
6. In the `hadoopEnv.properties` file, set the `HADOOP_CLASSPATH` to include the paths of the TDCH JAR files and the `hive-site.xml` file.

For example, set the `HADOOP_CLASSPATH` as follows:

```
infapdo.env.entry.hadoop_classpath=HADOOP_CLASSPATH=<Path of the dependent TDCH JAR files and the directory where the hive-site.xml file is stored>
```

7. In the `hadoopEnv.properties` file, set the `INFA_MAPRED_CLASSPATH` variable to include the paths of the TDCH JAR files and the `hive-site.xml` file.

For example, set the `INFA_MAPRED_CLASSPATH` as follows:

```
infapdo.env.entry.mapred_classpath=INFA_MAPRED_CLASSPATH=<Path of the dependent TDCH JAR files and the directory where the hive-site.xml file is stored>
```

Step 2. Configuring the InfaTDCHConfig.txt File

Configure the `InfaTDCHConfig.txt` file to define the parameters that TDCH must use to load data to Teradata in a Hive run-time environment. Informatica ships a sample `InfaTDCHConfig.txt` file. Customize the sample file based on your requirements.

1. Access the following directory: `<Informatica installation directory>\services\shared\hadoop\<Hadoop Distribution Directory>\infaConf`
2. Find the file named `InfaTDCHConfig.txt`.

Back up the file before you modify it.

Note: If the `InfaTDCHConfig.txt` file is missing from the `<Informatica installation directory>\services\shared\hadoop\biginsights_3.0.0.1\infaConf` directory, you must manually copy the file from another Hadoop Distribution Directory.

3. Use a text editor to open the file and edit the properties.
4. Set the `libjars` parameter to the paths that contain the TDCH JAR files.
5. Set other parameters based on how you want TDCH to process the data.
6. Save the file with the name `InfaTDCHConfig.txt`.

Step 3. Configuring the EnableTdch Custom Property

You must configure the EnableTdch custom property to enable TDCH to run a Teradata mapping in a Hive run-time environment.

1. Log in to the Informatica Administrator as an Administrator.
2. Select the Data Integration Service in the Domain Navigator to open the service properties.
3. In the contents panel, click the **Processes** view.
4. Click **Edit** in the **Custom Properties** section.
The **Edit Custom Properties** dialog box appears.
5. Click **New**.
The **New Custom Property** dialog box appears.
6. Enter the custom property name as EnableTdch and the value as True.
7. Click **OK** twice.
8. In the contents panel, click the **Properties** view.
9. Click **Recycle Service**.
The **Recycle Service** dialog box appears.
10. Choose to wait until all processes complete or to abort all processes.
11. Click **OK**.

The EnableTdch custom property is enabled at the Data Integration Service level.

CHAPTER 3

PowerExchange for Teradata Parallel Transporter Connections

This chapter includes the following topics:

- [Teradata Parallel Transporter Connection Overview, 16](#)
- [Teradata Parallel Transporter Connection Properties, 16](#)
- [Creating a Teradata Parallel Transporter Connection, 18](#)

Teradata Parallel Transporter Connection Overview

Use a Teradata PT connection to access Teradata tables.

Create a connection to import Teradata table metadata, create data objects, and run mappings. The Developer tool uses the connection when you import a data object. The Data Integration Service uses the connection when you run mappings.

Use the Developer tool, Administrator tool, or infacmd to create a Teradata PT connection.

Teradata Parallel Transporter Connection Properties

Use a Teradata PT connection to access Teradata tables. The Teradata PT connection is a database type connection. You can create and manage a Teradata PT connection in the Administrator tool or the Developer tool.

Note: The order of the connection properties might vary depending on the tool where you view them.

The following table describes Teradata PT connection properties:

Property	Description
Name	Name of the connection. The name is not case sensitive and must be unique within the domain. The name cannot exceed 128 characters, contain spaces, or contain the following special characters: ~ ` ! \$ % ^ & * () - + = { [] } \ : ; " ' < , > . ? /
ID	String that the Data Integration Service uses to identify the connection. The ID is not case sensitive. It must be 255 characters or less and must be unique in the domain. You cannot change this property after you create the connection. Default value is the connection name.
Description	Description of the connection. The description cannot exceed 765 characters.
User Name	Teradata database user name with the appropriate write permissions to access the database.
Password	Password for the Teradata database user name.
Driver Name	Name of the Teradata JDBC driver.
Connection String	Connection string used to access metadata from the database. Use the following connection string: <code>jdbc:teradata://<hostname>/database=<database name>, tmode=ANSI, charset=UTF8</code>

The following table describes the properties for data access:

Property	Description
TDPID	Name or IP address of the Teradata database machine.
Database Name	Teradata database name. If you do not enter a database name, Teradata PT API uses the default login database name.
Data Code Page	Code page associated with the database. When you run a mapping that loads to a Teradata target, the code page of the Teradata PT connection must be the same as the code page of the Teradata target. Default is UTF-8.
Tenacity	Number of hours that Teradata PT API continues trying to log on when the maximum number of operations run on the Teradata database. Must be a positive, non-zero integer. Default is 4.
Max Sessions	Maximum number of sessions that Teradata PT API establishes with the Teradata database. Must be a positive, non-zero integer. Default is 4.

Property	Description
Min Sessions	Minimum number of Teradata PT API sessions required for the Teradata PT API job to continue. Must be a positive integer between 1 and the Max Sessions value. Default is 1.
Sleep	Number of minutes that Teradata PT API pauses before it retries to log on when the maximum number of operations run on the Teradata database. Must be a positive, non-zero integer. Default is 6.
Use Metadata JDBC URL for TDCH	Indicates that the Teradata Connector for Hadoop (TDCH) must use the JDBC URL that you specified in the connection string under the metadata access properties. Default is selected. Clear this option to enter a different JDBC URL that TDCH must use when it runs the mapping.
Data Encryption	Enables full security encryption of SQL requests, responses, and data on Windows. To enable data encryption on Unix, add the command <code>UseDataEncryption=Yes</code> to the DSN in the <code>odbc.ini</code> file. Default is disabled.
TDCH JDBC Url	Enter the JDBC URL that TDCH must use when it runs a Teradata mapping. Use the following format: <code>jdbc:teradata://<hostname>/database=<database name>,tmode=ANSI,charset=UTF8</code> This field is available only when you clear the Use Metadata JDBC URL for TDCH option.

Creating a Teradata Parallel Transporter Connection

Create a Teradata PT connection in the Developer tool before you import Teradata data objects or run mappings.

1. Click **Window > Preferences**.
2. Click **Informatica > Connections**.
3. Expand the domain.
4. Select **Databases > Teradata PT** and click **Add**.
5. Enter a connection name.
6. Enter an ID for the connection.
7. Optionally, enter a connection description.
8. Select the domain where you want to create the connection.
9. Select the type of connection.
10. Click **Next**.
11. Configure the connection properties for metadata access and data access.

12. Click **Test Connection** to verify that you can connect to the Teradata database.
The properties for data access are validated at run time.
13. Click **Finish**.

CHAPTER 4

PowerExchange for Teradata Parallel Transporter API Data Objects

This chapter includes the following topics:

- [Teradata Data Object Overview, 20](#)
- [Teradata Data Object Sections, 21](#)
- [Teradata Data Object Overview Properties, 21](#)
- [Teradata Data Object Read Operation Properties, 22](#)
- [Teradata Data Object Write Operation Properties, 24](#)
- [Importing a Teradata Data Object, 30](#)
- [Creating a Teradata Data Object Read Operation, 31](#)
- [Creating a Teradata Data Object Write Operation, 31](#)

Teradata Data Object Overview

A Teradata data object is a physical data object that uses Teradata as a source and a target. A Teradata data object is the representation of data that is based on a Teradata table. You can configure the data object read and write operation properties that determine how data can be read from or loaded to Teradata tables.

Import the Teradata table into the Developer tool to create a Teradata data object. Create a data object read or write operation for the Teradata data object. Then, you can add the data object read or write operation to a mapping.

Teradata Data Object Sections

The Teradata data object contains sections to edit the object name and the properties.

After you create a Teradata data object, you can change the data object properties in the following data object views:

- **Overview** view. Edit the Teradata data object name, description, and tables.
- **Data Object Operation** view. Edit the properties that the Data Integration Service uses when it reads data from or writes data to a Teradata table.

When you create a mapping for a Teradata read or write operation, you can view the data object properties in the **Properties** section.

Teradata Data Object Overview Properties

The Teradata **Overview** section displays general information about the Teradata data object and detailed information about the Teradata table that you imported.

The following table describes the general properties that you configure for a Teradata data object:

Property	Description
Name	Name of the Teradata data object.
Description	Description of the Teradata data object.
Connection	Name of the Teradata PT connection. Click Browse to select another Teradata PT connection.

The following table describes the properties of the Teradata table that you import:

Property	Description
Name	Name of the Teradata table.
Type	Native data type of the Teradata table.
Description	Description of the Teradata table.

Teradata Data Object Read Operation Properties

The Data Integration Service reads data from a Teradata table based on the data object read operation properties. The Developer tool displays the data object read operation properties for the Teradata data object in the **Data Object Operation** section.

You can view or configure the data object read operation from the input and target properties.

- **Input properties.** Represents data that the Data Integration Service passes into the mapping pipeline. Select the input properties to edit the port properties and specify the advanced properties of the data object read operation.
- **Source properties.** Represents data that the Data Integration Service reads from the Teradata table. Select the source properties to view data such as the name and description of the Teradata table. Create key columns to identify each row in a data object read operation, and create key relationships between data object read operations.

Input Properties of the Data Object Read Operation

The input properties represent data that the Data Integration Service passes into the mapping pipeline. Select the input properties to edit the port properties of the data object read operation. You can also specify advanced data object read operation properties to read data from Teradata tables.

The input properties of the data object read operation include general properties that apply to the data object read operation. They also include port, source, and advanced properties that apply to the data object read operation.

You can edit the input properties of the data object read operation from the **Ports** and **Advanced** tabs.

Sources Properties

The sources properties list the Teradata tables in the data object read operation.

Advanced Properties

The advanced properties allow you to specify data object read operation properties to import data from Teradata tables.

The following table describes the advanced properties that you configure in the data object read operation:

Property	Description
Pre SQL	SQL command that the Data Integration Service runs against the source database before it reads from the source.
Post SQL	SQL command that the Data Integration Service runs against the source database before it writes to the target.
Select Distinct	Specifies if you want to select only unique rows. The Data Integration Service includes a SELECT DISTINCT statement if you choose this option.
Select SQL Statement	Specifies the columns that you want to use from the source database.

Property	Description
Query Band Expression	<p>The query band expression to be passed to the Teradata PT API.</p> <p>A query band expression is a set of name-value pairs that identify a query's originating source. In the expression, each name-value pair is separated by a semicolon and the expression ends with a semicolon. For example, <code>ApplicationName=Informatica;Version=9.5.0;ClientUser=A;</code>.</p>
Source Table Name	Table name from which you want to extract rows.
Source Table Schema	Schema name from which you want to extract rows.
Driver Tracing Level	<p>Determines Teradata PT API tracing at the driver level:</p> <ul style="list-style-type: none"> - TD_OFF. Teradata PT API disables tracing. - TD_OPER. Teradata PT API enables tracing for driver-specific activities for Teradata. - TD_OPER_ALL. Teradata PT API enables all driver-level tracing. - TD_OPER_CLI. Teradata PT API enables tracing for activities involving CLIV2. - TD_OPER_NOTIFY. Teradata PT API enables tracing for activities involving the Notify feature. - TD_OPER_OPCOMMON. Teradata PT API enables tracing for activities involving the operator common library. <p>Default is TD_OFF.</p> <p>The input locale for the value of the Driver Tracing Level attribute is in the English language only.</p>
Infrastructure Tracing Level	<p>Determines Teradata PT API tracing at the infrastructure level.</p> <p>Select one of the following values:</p> <ul style="list-style-type: none"> - TD_OFF. Teradata PT API disables tracing. - TD_OPER. Teradata PT API enables tracing for driver-specific activities for Teradata. - TD_OPER_ALL. Teradata PT API enables all driver-level tracing. - TD_OPER_CLI. Teradata PT API enables tracing for activities involving CLIV2. - TD_OPER_NOTIFY. Teradata PT API enables tracing for activities involving the Notify feature. - TD_OPER_OPCOMMON. Teradata PT API enables tracing for activities involving the operator common library. <p>Default is TD_OFF.</p> <p>You must enable the driver tracing level before you can enable the infrastructure tracing level.</p> <p>The input locale for the value of the Infrastructure Tracing Level attribute is in the English language only.</p>
Trace File Name	<p>File name and path of the Teradata PT API trace file.</p> <p>Default path is <code>\$(Informatica installation directory)\tomcat\bin</code>.</p> <p>Default file name is <code><Name of the TPT Operator>_timestamp</code>. For example, <code>LOAD_20091221</code>.</p>
Spool Mode	<p>Determines the spool mode Teradata PT API uses to extract data from Teradata.</p> <p>You can choose one of the following spool modes:</p> <ul style="list-style-type: none"> - Spool. Teradata PT API spools data while extracting data from Teradata. - NoSpool. Teradata PT API does not spool data while extracting data from Teradata. If the database does not support the NoSpool option, Teradata PT API uses the Spool option. - NoSpoolOnly. Teradata PT API does not spool while extracting data from Teradata. <p>Default is Spool.</p> <p>The input locale for the value of the Spool Mode attribute is in the English language only.</p>

Source Properties of the Data Object Read Operation

The source properties represent the data that is populated based on the Teradata tables that you added when you created the data object. The source properties of the data object read operation include general and column properties that apply to the Teradata tables.

You can edit the source properties of the data object read operation from the **Keys** and **Column** tabs.

Keys Properties

Displays the Teradata table that has a primary key.

You can view the name and description of the table that has a primary key. The **Available Fields** section displays the columns that are not part of the primary key. The **Selected Fields** section displays the columns that are part of the primary key.

Column Properties

The column properties display the data types, precision, and scale of the target property in the data object read operation.

The following table describes the column properties of the data object read operation:

Property	Description
Name	Name of the column property.
Type	Native data type of the column property.
Precision	Maximum number of significant digits for numeric data types, or maximum number of characters for string data types. For numeric data types, precision includes scale.
Scale	Maximum number of digits after the decimal point for numeric values.
Primary Key	Determines if the column property is a part of the primary key.
Description	Description of the column property.

Teradata Data Object Write Operation Properties

The Data Integration Service writes data to a Teradata table based on the data object write operation properties. The Developer tool displays the data object write operation properties for the Teradata data object in the **Data Object Operation** section.

You can view or configure the data object write operation from the input and target properties.

- Input properties. Represents data that the Data Integration Service passes into the mapping pipeline. Select the input properties to edit the port properties and specify the advanced properties of the data object write operation.

- **Target properties.** Represents data that the Data Integration Service writes to the Teradata table. Select the target properties to view data such as the name and description of the Teradata table, create key columns to identify each row in a data object write operation, and create key relationships between data object write operations.

Input Properties of the Data Object Write Operation

The input properties represent data that the Data Integration Service passes into the mapping pipeline. Select the input properties to edit the port properties of the data object write operation. You can also specify advanced data object write operation properties to load data into Teradata tables.

The input properties of the data object write operation include general properties that apply to the data object write operation. They also include port, source, and advanced properties that apply to the data object write operation.

You can edit the input properties of the data object write operation from the **Ports** and **Advanced** tabs.

Ports Properties

The input ports properties list the data types, precision, and scale of the data object write operation.

The following table describes the input ports properties that you configure in a data object write operation:

Property	Description
Name	Name of the port property.
Type	Data type of the port property.
Precision	Maximum number of significant digits for numeric data types, or maximum number of characters for string data types. For numeric data types, precision includes scale.
Scale	Maximum number of digits after the decimal point for numeric values.
Description	Description of the port property.

Advanced Properties

The advanced properties allow you to specify data object write operation properties to load data into Teradata tables.

The following table describes the advanced properties that you configure in the data object write operation:

Property	Description
Macro Database	<p>Name of the database that stores the macros Teradata PT API creates when you select the Stream system operator.</p> <p>The Stream system operator uses macros to modify tables. It creates macros before Teradata PT API begins loading data and removes them from the database after Teradata PT API loads all rows to the target.</p> <p>If you do not specify a macro database, Teradata PT API stores the macros in the log database.</p>
Truncate Table	<p>Teradata PT API deletes all rows in the Teradata target before it loads data.</p> <p>This attribute is available for the Stream, Update, and Load system operators.</p> <p>Default is disabled.</p>
Pause Acquisition	<p>Causes the load operation to pause before writing data to the Teradata PT API target.</p> <p>Default is disabled.</p>
Query Band Expression	<p>Query band expression to be passed to Teradata PT API.</p> <p>A query band expression is a set of name-value pairs that identify a query's originating source. In the expression, each name-value pair is separated by a semicolon and the expression ends with a semicolon. For example, <code>ApplicationName=Informatica;Version=9.5.0;ClientUser=A;</code>.</p>
Mark Duplicate Rows	<p>Specifies how Teradata PT API handles duplicate rows when it tries to insert or upsert rows in the target table.</p> <p>Select one of the following values:</p> <ul style="list-style-type: none">- None. If Teradata PT API receives a row marked for insert or upsert that causes a duplicate row in the target table, Teradata PT API does not write the row to the error table and does not mark it as an error row in the workflow log.- For Insert. If Teradata PT API receives a row marked for insert that exists in the target table, Teradata PT API writes the row to the error table and marks it as an error row in the workflow log.- For Update. If Teradata PT API receives a row marked for update that causes a duplicate row in the target table, Teradata PT API writes the row to the error table and marks it as an error row in the workflow log.- Both. If Teradata PT API receives a row marked for insert or upsert that causes a duplicate row in the target table, Teradata PT API writes the row to the error table and marks it as an error row in the workflow log. <p>Default is For Insert.</p> <p>The input locale for the value of the Mark Duplicate Rows attribute is in the English language only.</p>
Log Database	<p>Name of the database that stores the log tables.</p>
Log Table Name	<p>Name of the restart log table.</p>
Error Database	<p>Name of the database that stores the error tables.</p>
Error TableName1	<p>Name of the first error table. Use this attribute to override the default error table name.</p>

Property	Description
Error TableName2	Name of the second error table. Use this attribute to override the default error table name.
Drop Log/Error	Drops the existing log and work tables when you run a mapping. This attribute is available for a Teradata target ODBC connection. Default is disabled.
Serialize	Uses the Teradata PT API serialize mechanism to reduce locking overhead when you select the Stream system operator. Default is enabled.
Pack	Number of statements to pack into a request when you select the Stream system operator. Must be a positive, nonzero integer. Default is 20. Minimum is 1. Maximum is 600.
Pack Maximum	Causes Teradata PT API to determine the maximum number of statements to pack into a request when you select the Stream system operator. Default is disabled.
Buffers	Determines the maximum number of request buffers that may be allocated for the Teradata PT API job when you select the Stream system operator. Must be a positive, nonzero integer. Default is 3. Minimum is 2.
Error Limit	Maximum number of records that can be stored in the error table before Teradata PT API terminates the Stream system operator job. Must be -1 or a positive, nonzero integer. Default is -1, which specifies an unlimited number of records.
Replication Override	Specifies how Teradata PT API overrides the normal replication services controls for an active Teradata PT API mapping. Select one of the following values: <ul style="list-style-type: none"> - On. Teradata PT API overrides normal replication services controls for the active mapping. - Off. Teradata PT API disables override of normal replication services for the active mapping when change data capture is active. - None. Teradata PT API does not send an override request to the Teradata Database. Default is None. The input locale for the value of the Replication Override attribute is in the English language only.
Driver Tracing Level	Determines Teradata PT API tracing at the driver level. Select one of the following values: <ul style="list-style-type: none"> - TD_OFF. Teradata PT API disables tracing. - TD_OPER. Teradata PT API enables tracing for driver-specific activities for Teradata. - TD_OPER_ALL. Teradata PT API enables all driver-level tracing. - TD_OPER_CLI. Teradata PT API enables tracing for activities involving CLlV2. - TD_OPER_NOTIFY. Teradata PT API enables tracing for activities involving the Notify feature. - TD_OPER_OPCOMMON. Teradata PT API enables tracing for activities involving the operator common library. Default is TD_OFF. The input locale for the value of the Driver Tracing Level attribute is in the English language only.

Property	Description
Infrastructure Tracing Level	<p>Determines Teradata PT API tracing at the infrastructure level.</p> <p>Select one of the following values:</p> <ul style="list-style-type: none"> - TD_OFF. Teradata PT API disables tracing. - TD_OPER. Teradata PT API enables tracing for driver-specific activities for Teradata. - TD_OPER_ALL. Teradata PT API enables all driver-level tracing. - TD_OPER_CLI. Teradata PT API enables tracing for activities involving CLlV2. - TD_OPER_NOTIFY. Teradata PT API enables tracing for activities involving the Notify feature. - TD_OPER_OPCOMMON. Teradata PT API enables tracing for activities involving the operator common library. <p>Default is TD_OFF.</p> <p>You must enable the driver tracing level before you can enable the infrastructure tracing level.</p> <p>The input locale for the value of the Infrastructure Tracing Level attribute is in the English language only.</p>
Trace File Name	<p>File name and path of the Teradata PT API trace file.</p> <p>Default path is <code>\$<Informatica installation directory>\tomcat\bin</code>.</p> <p>Default file name is <code><Name of the TPT Operator>_timestamp</code>. For example, <code>LOAD_20091221</code>.</p>
System Operator	<p>Specifies the Teradata PT API operator type.</p> <p>Select one of the following values:</p> <ul style="list-style-type: none"> - Load. Performs bulk loading into an empty Teradata database table. You can use the Load system operator in native and Hive environments. - Stream. Performs insert, update, upsert, and delete operations against Teradata database tables in near real-time mode. You can use the Stream system operator in native and Hive environments. - Update. Performs insert, update, upsert, and delete operations against Teradata database tables. You can use the Update system operator in native and Hive environments. <p>Default is Stream.</p> <p>The input locale for the value of the System Operator attribute is in the English language only.</p>
Mark Missing Rows	<p>Specifies how Teradata PT API handles rows that do not exist in the target table when it tries to update or delete rows.</p> <p>Select one of the following values:</p> <ul style="list-style-type: none"> - None. If Teradata PT API receives a row marked for update or delete but it is missing in the target table, Teradata PT API does not write the row to the error table and does not mark it as an error row in the workflow log. - For Update. If Teradata PT API receives a row marked for update but it is missing in the target table, Teradata PT API writes the row to the error table and marks it as an error row in the workflow log. - For Delete. If Teradata PT API receives a row marked for delete but it is missing in the target table, Teradata PT API writes the row to the error table and marks it as an error row in the workflow log. - Both. If Teradata PT API receives a row marked for update or delete but it is missing in the target table, Teradata PT API writes the row to the error table and marks it as an error row in the workflow log. <p>Default is None.</p> <p>The input locale for the value of the Mark Missing Rows attribute is in the English language only.</p>

Property	Description
Mark Extra Rows	<p>Specifies how Teradata PT API marks error rows when it tries to update or delete multiple target table rows.</p> <p>Select one of the following values:</p> <ul style="list-style-type: none"> - None. If Teradata PT API receives a row marked for update or delete that affects multiple rows in the target table, Teradata PT API does not write the row to the error table and does not mark it as an error row in the workflow log. - For Update. If Teradata PT API receives a row marked for update that affects multiple rows in the target table, Teradata PT API writes the row to the error table and marks it as an error row in the workflow log. - For Delete. If Teradata PT API receives a row marked for delete that affects multiple rows in the target table, Teradata PT API writes the row to the error table and marks it as an error row in the workflow log. - Both. If Teradata PT API receives a row marked for update or delete that affects multiple rows in the target table, Teradata PT API writes the row to the error table and marks it as an error row in the workflow log. <p>Default is Both.</p> <p>This attribute is available only for the Stream system operator.</p> <p>The input locale for the value of the Mark Extra Rows attribute is in the English language only.</p>
Insert	<p>Teradata PT API creates a DML group to insert rows.</p> <p>If you do not want to insert rows, clear this option to increase session performance.</p> <p>Default is selected.</p>
Update	<p>Teradata PT API creates a DML group to update rows.</p> <p>Default is not selected.</p>
Upsert	<p>Teradata PT API creates a DML group to update or insert rows.</p> <p>Default is not selected.</p>
Delete	<p>Teradata PT API creates a DML group to delete rows.</p> <p>Default is not selected.</p>
WorkTable Database Name	<p>Name of the database that stores the work tables.</p> <p>Use this attribute to override the default work table database name. If you do not specify a database name, the Data Integration Service uses the target table database.</p>
WorkTable Name	<p>Name of the work tables when you select the Update system operator.</p> <p>The Teradata database creates one work table for each target table. If you do not specify a work table name, the Data Integration Service uses the name <work_table_database>.INFA<number>_WT. The exact table name appears in the session log.</p>
Pre SQL	<p>Pre-session SQL commands to run against the source database before the Data Integration Service reads the source.</p>
Post SQL	<p>Post-session SQL commands to run against the source database after the Data Integration Service writes to the target.</p>

Target Properties of the Data Object Write Operation

The target properties represent the data that is populated based on the Teradata tables that you added when you created the data object. The target properties of the data object write operation include general and column properties that apply to the Teradata tables.

You can edit the target properties of the data object write operation from the **Keys** and **Column** tabs.

Keys Properties

Displays the Teradata table that has a primary key.

You can view the name and description of the table that has a primary key. The **Available Fields** section displays the columns that are not part of the primary key. The **Selected Fields** section displays the columns that are part of the primary key.

Column Properties

The column properties display the data types, precision, and scale of the target property in the data object write operation.

The following table describes the target column properties of the data object write operation:

Property	Description
Name	Name of the column property.
Type	Native data type of the column property.
Precision	Maximum number of significant digits for numeric data types, or maximum number of characters for string data types. For numeric data types, precision includes scale.
Scale	Maximum number of digits after the decimal point for numeric values.
Primary Key	Determines if the column property is a part of the primary key.
Description	Description of the column property.

Importing a Teradata Data Object

Import a Teradata data object to add to a mapping.

Configure a Teradata PT connection, before you import a Teradata data object.

1. Select a project or folder in the **Object Explorer** view.
2. Click **File > New > Data Object**.
3. Select **Teradata Data Object** and click **Next**.
The **New Teradata Data Object** dialog box appears.
4. Enter a name for the data object.
5. Click **Browse** next to the **Location** option and select the target project or folder.
6. Click **Browse** next to the **Connection** option and select the Teradata PT connection from which you want to import the Teradata table metadata.
7. To add a table, click **Add** next to the **Selected Resources** option.
The **Add Resource** dialog box appears.
8. Select a table. You can search for it or navigate to it.
 - Navigate to the Teradata table that you want to import and click **OK**.

- To search for a table, enter the name or the description of the table you want to add. Click **OK**.
9. If required, add more tables to the Teradata data object.
You can also add tables to a Teradata data object after you create it.
 10. Click **Finish**.
The data object appears under Data Objects in the project or folder in the **Object Explorer** view.

Creating a Teradata Data Object Read Operation

You can create the data object read operation for one or more Teradata data objects. You can add a Teradata data object read operation to a mapping as a source.

Before you create a Teradata data object read operation, you must create a Teradata data object with at least one Teradata table.

1. Select the data object in the Object Explorer view.
2. Right-click and select **New > Data Object Operation**.
The **Data Object Operation** dialog box appears.
3. Enter a name for the data object operation.
4. Select **Read** as the type of data object operation.
5. Click **Add**.
The **Select Resources** dialog box appears.
6. Select the Teradata table for which you want to create the data object read operation and click **OK**.
You can select only one data object to a data object read operation.
7. Click **Finish**.
The Developer tool creates the data object read operation for the selected data object.

Creating a Teradata Data Object Write Operation

You can create the data object write operation for one or more Teradata data objects. You can add a Teradata data object write operation to a mapping as a target.

Before you create a Teradata data object write operation, you must create the Teradata data object with at least one Teradata table.

1. Select the data object in the Object Explorer view.
2. Right-click and select **New > Data Object Operation**.
The **Data Object Operation** dialog box appears.
3. Enter a name for the data object operation.
4. Select **Write** as the type of data object operation.
5. Click **Add**.
The **Select Resources** dialog box appears.

6. Select the Teradata table for which you want to create the data object write operation and click **OK**.
You can select only one data object to a data object write operation.
7. Click **Finish**.

The Developer tool creates the data object write operation for the selected data object.

CHAPTER 5

PowerExchange for Teradata Parallel Transporter API Mappings

This chapter includes the following topics:

- [Teradata Mappings Overview, 33](#)
- [Hive Validation and Run-time Environments, 33](#)
- [Teradata Mapping Example, 34](#)

Teradata Mappings Overview

After you create a Teradata data object read or write operation, you can create a mapping.

You can create an Informatica mapping containing a Teradata data object read operation as the input, and a relational or flat file data object operation as the target. You can create an Informatica mapping containing objects such as a relational or flat file data object operation as the input, transformations, and a Teradata data object write operation as the output to load data to Teradata tables.

Validate and run the mapping. You can also add the mapping to a Mapping task in a workflow and run the workflow.

Hive Validation and Run-time Environments

You configure the mappings to run in native or Hive run-time environment. When you run mappings in the native environment, the Data Integration Service processes the mapping. When you run a mapping in a Hive environment, the Data Integration Service can push mappings that are developed in the Developer tool to a Hadoop cluster.

For more information about the Hive run-time and validation environment, see the *Informatica Big Data Management User Guide*.

Teradata Mapping Example

Your organization, ABC Airlines, needs to analyze flight information from the log files that are generated for each flight to find the number of flights to a particular destination in a month. Create a mapping that reads the flight log information of all the flights and writes specific information to Teradata tables.

You can use the following objects in a Teradata mapping:

Mapping input

The mapping source is a table in an Oracle database or a flat file that contains the flight log information of ABC Airlines.

Import the Oracle or flat file data object with the flight log information. Create a data object read operation for the Oracle data object you have created. Add the data object read operation to the mapping.

Transformation

Use the Filter transformation to specify the month and destination.

Mapping output

The mapping output is a data object write operation to load the data to a Teradata table. Run the mapping either in native or Hive environment.

When you run the mapping, the extracted information about the number of flights in a month to a specific destination is written to the Teradata target table. You can then use the information for data analysis.

CHAPTER 6

Teradata Parallel Transporter API Run-time Processing

This chapter includes the following topics:

- [Teradata Run-time Processing Overview, 35](#)
- [System Operators, 35](#)
- [Error and Log Tables, 36](#)
- [Parameterization, 39](#)
- [Partitioning, 40](#)

Teradata Run-time Processing Overview

When you create a Teradata mapping, you define data object operation read or write properties that determine how data is read from Teradata tables or loaded to Teradata targets.

You can set the following data object operation read or write properties:

- System operator
- Error limits
- Names of error tables and log tables
- Strategy that Teradata PT API must use to treat duplicate rows, missing rows, and extra rows when it loads data to a Teradata target

System Operators

The Data Integration Service uses Teradata PT API to connect to Teradata. It loads data by using the Teradata PT API Stream, Load, or Update system operators.

Stream System Operator

You can use the Stream system operator to perform high-speed parallel inserts to empty or existing Teradata tables without locking target tables. You can also perform update, upsert, and delete operations after you define a primary key in the target table.

You can maintain current and accurate data for immediate decision making. The Stream system operator uses macros to modify tables. The Stream system operator creates macros before Teradata PT API begins loading data and removes them from the database after Teradata PT API loads all rows to the target. If you do not specify a macro database, Teradata PT API stores the macros in the log database. You can use the Stream system operator in native and Hive environments.

Error Limit for the Stream System Operator

Teradata PT API terminates the Stream system operator job after it reaches the maximum number of records that can be stored in the error table.

The error limit is approximate because the Stream operator sends multiple rows of data at a time to the Teradata table. By the time Teradata PT API processes the message that indicates that the error limit has been exceeded, it might have loaded more records into the error table than the number specified in the error limit.

By default, the error limit value is unlimited. The error limit applies to each instance of the Stream system operator.

Load System Operator

The Load system operator loads a large volume of data at high speed into an empty Teradata table. The Load operator is often used to load initial data to Teradata tables as it inserts data into individual rows of a target table. You can use the Load system operator in native and Hive environments.

Update System Operator

You can use the Update system operator to perform insert, update, upsert, and delete operations against Teradata database tables. You can load data to empty or existing Teradata tables.

If you want to run an update, upsert, or delete operation on a Teradata target table, you must ensure that the target has a primary key column.

You can use the Update system operator in native and Hive environments.

Error and Log Tables

Use the log and error tables to view log, error, and rejected data when a mapping runs.

When you use Teradata PT API to run a mapping that loads data to Teradata, Teradata PT API creates log and error tables. The log tables store Teradata PT API restart and log information. The error tables log Teradata errors and rejected data when a mapping runs.

Error Tables

Teradata writes rejected data to error tables ErrorTable1 and ErrorTable2.

ErrorTable1 contains data that was rejected for the following reasons:

- Data conversion errors
- Constraint violations
- Access Module Processor configuration changes

ErrorTable2 contains data that was rejected for the following reasons:

- Unique primary index constraint violations
- Load driver job acquisition phase errors

When you configure a Teradata mapping, you can enter a name for each error table. You can create the error tables in any database. By default, the error tables are stored in the default database. Choose where you want to create the error tables when you configure a mapping to load to Teradata.

When a mapping fails, see the error tables for more information about the errors. Before you run the mapping again, drop the error tables or enter different table names in the data object operation write properties. Otherwise, the mapping fails.

The following table describes the data object operation write properties that allow you to specify error table names:

Property	Description
Error Database	Name of the database that stores the error tables. If you do not enter an error database name in the data object operation write properties, Teradata PT API stores the error tables in the database you specify in the connection object.
Error Table Name1	Name of the first error table. If you do not specify a name for the first error table, the Data Integration Service uses the name <code><error_database>.INFA_ET1_<number></code> . The table name appears in the mapping log.
Error Table Name2	Name of the second error table. If you do not specify a name for the second error table, the Data Integration Service uses the name <code><error_database>.INFA_ET2_<number></code> . The table name appears in the mapping log.

Log Tables

When you configure a Teradata mapping, you can enter a log table name. You can create the error tables in any database. By default, the log tables are stored in the default database.

When a mapping fails, see the log table for more information. Before you run the mapping again, drop the log table or enter a different table name in the data object operation write properties. Otherwise, the mapping fails.

The following table describes the data object operation write properties that allow you to specify log table information:

Property	Description
Log Database	Name of the database that stores the log tables. If you do not enter a log database name in the data object operation write properties, Teradata PT API stores the log tables in the database you specify in the connection object.
Log Table Name	Name of the log table. If you do not specify a log table name, the Data Integration Service uses the name <code><log_database>.INFA_LT_<number></code> . The table name appears in the mapping log.

Row-level Error Processing

When you write data to a Teradata target, you can configure how Teradata PT API treats duplicate rows, missing rows, and extra rows.

Duplicate Rows

You can configure how Teradata PT API handles duplicate rows when it tries to insert or upsert rows in the target table.

You can select one of the following values:

- None. If Teradata PT API receives a row marked for insert or upsert that causes a duplicate row in the target table, Teradata PT API does not write the row to the error table and does not mark it as an error row in the workflow log.
- For Insert. If Teradata PT API receives a row marked for insert that exists in the target table, Teradata PT API writes the row to the error table and marks it as an error row in the workflow log.
- For Update. If Teradata PT API receives a row marked for update that causes a duplicate row in the target table, Teradata PT API writes the row to the error table and marks it as an error row in the workflow log.
- Both. If Teradata PT API receives a row marked for insert or upsert that causes a duplicate row in the target table, Teradata PT API writes the row to the error table and marks it as an error row in the workflow log.

Missing Rows

You can configure how Teradata PT API handles rows that do not exist in the target table when it tries to update or delete rows.

You can select one of the following values:

- None. If Teradata PT API receives a row marked for update or delete but it is missing in the target table, Teradata PT API does not write the row to the error table and does not mark it as an error row in the workflow log.
- For Update. If Teradata PT API receives a row marked for update but it is missing in the target table, Teradata PT API writes the row to the error table and marks it as an error row in the workflow log.
- For Delete. If Teradata PT API receives a row marked for delete but it is missing in the target table, Teradata PT API writes the row to the error table and marks it as an error row in the workflow log.
- Both. If Teradata PT API receives a row marked for update or delete but it is missing in the target table, Teradata PT API writes the row to the error table and marks it as an error row in the workflow log.

Extra Rows

You can configure how Teradata PT API marks error rows when it tries to update or delete multiple target table rows.

You can select one of the following values:

- None. If Teradata PT API receives a row marked for update or delete that affects multiple rows in the target table, Teradata PT API does not write the row to the error table and does not mark it as an error row in the workflow log.
- For Update. If Teradata PT API receives a row marked for update that affects multiple rows in the target table, Teradata PT API writes the row to the error table and marks it as an error row in the workflow log.
- For Delete. If Teradata PT API receives a row marked for delete that affects multiple rows in the target table, Teradata PT API writes the row to the error table and marks it as an error row in the workflow log.
- Both. If Teradata PT API receives a row marked for update or delete that affects multiple rows in the target table, Teradata PT API writes the row to the error table and marks it as an error row in the workflow log.

This attribute is available only for the Stream system operator.

Parameterization

You can parameterize Teradata data object operation properties to override the read and write data object operation properties during run time.

In bulk mode, you can parameterize the following data object read operation properties for Teradata sources:

- Pre SQL
- Post SQL
- Select Distinct
- Select SQL Statement
- Query Band Expression
- Source Table Name
- Source Table Schema
- Driver Tracing Level
- Infrastructure Tracing Level
- Trace File Name
- Spool Mode

In bulk mode, you can parameterize the following data object write operation properties for Teradata targets:

- Macro Database
- Truncate Table
- Pause Acquisition
- Query Band Expression
- Mark Duplicate Rows
- Log Database

- Log Table Name
- Error Database
- Error Table Name1
- Error Table Name2
- Drop Log/Error
- Serialize
- Pack
- Pack Minumum
- Buffers
- Error Limit
- Replication Override
- Driver Tracing Level
- Infrastructure Tracing Level
- Trace File Name
- System Operator
- Mark MISSING Rows
- Mark EXTRA Rows
- Insert
- Update
- Upsert
- Delete
- WorkTable Database Name
- WorkTable Name
- Pre SQL
- Post SQL

The following attributes support partial parameterization:

- Query Band Expression
- Trace File Name

Partitioning

When a mapping that is enabled for partitioning contains a Teradata data object as a target, the Data Integration Service can use multiple threads to write to the target.

You can configure dynamic partitioning for Teradata Parallel Transporter API data objects. You can configure the partition information so that the Data Integration Service determines the number of partitions to create at run time.

To enable partitioning, administrators and developers perform the following tasks:

Administrators set maximum parallelism for the Data Integration Service to a value greater than 1 in the Administrator tool.

Maximum parallelism determines the maximum number of parallel threads that process a single pipeline stage. Administrators increase the **Maximum Parallelism** property value based on the number of CPUs available on the nodes where mappings run.

Developers can set a maximum parallelism value for a mapping in the Developer tool.

By default, the **Maximum Parallelism** property for each mapping is set to Auto. Each mapping uses the maximum parallelism value defined for the Data Integration Service.

Developers can change the maximum parallelism value in the mapping run-time properties to define a maximum value for a particular mapping. When maximum parallelism is set to different integer values for the Data Integration Service and the mapping, the Data Integration Service uses the minimum value of the two.

Note: When you create a Teradata write data object, ensure that the **Serialize** option is set to **Off** to enable partitioning for the Stream operator.

Dynamic Partitioning

If the volume of data grows or you add more processors, you can adjust partitioning so that the run time does not increase. When you use dynamic partitioning, you can configure the partition information so that the Data Integration Service determines the number of partitions to create at run time.

Some transformations do not support partitioning. When a mapping enabled for partitioning contains a transformation that does not support partitioning, the Data Integration Service uses one thread to run the transformation. The Data Integration Service can use multiple threads to run the remaining mapping pipeline stages.

You can optimize write performance by using dynamic partitioning.

APPENDIX A

Data Type Reference

This appendix includes the following topics:

- [Data Type Reference Overview, 42](#)
- [Teradata Data Types and Transformation Data Types, 42](#)

Data Type Reference Overview

Informatica Developer uses the following data types in Teradata mappings:

- Teradata native data types. Teradata data types appear in the data object operation column properties.
- 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 Data Integration Service uses to move data across platforms. Transformation data types appear in all transformations in a mapping.

When the Data Integration Service writes to a target, it converts the transformation data types to the comparable native data types.

Teradata Data Types and Transformation Data Types

The following table lists the Teradata data types that the Data Integration Service supports and the corresponding transformation data types:

Teradata Data Type	Range	Transformation Data Type	Range
Bigint	-9,223,372,036,854,775,808 to 9,223,372,036,854,775,807	Bigint	-9,223,372,036,854,775,808 to 9,223,372,036,854,775,807 Precision 19, scale 0
Byte	1 to 64,000 bytes	Binary	1 to 104,857,600 bytes
Byteint	-128 to 127	Small Integer	Precision 5, scale 0
Char	1 to 64,000 bytes	String	1 to 104,857,600 characters

Teradata Data Type	Range	Transformation Data Type	Range
Date	Jan 1, 0001 A.D. to Dec 31, 9999 A.D. Precision 19, scale 0	Date/Time	Jan 1, 0001 A.D. to Dec 31, 9999 A.D. (precision to the nanosecond)
Decimal	Precision 1 to 18, scale 0 to 18	Decimal	Precision 1 to 28, scale 0 to 28
Float	-2.226E+308 to 1.797E+308	Double	Precision 15
Integer	-2,147,483,648 to 2,147,483,647	Integer	-2,147,483,648 to 2,147,483,647 Precision 10, scale 0
Smallint	-32768 to 32768	Small Integer	Precision 5, scale 0
Time	00:00:00.000000 to 23:59:61.999999 Precision 15, scale 6	Date/Time	Jan 1, 0001 A.D. to Dec 31, 9999 A.D. (precision to the nanosecond)
Timestamp	1 to 19 characters Precision 19 to 26, scale 0 to 6	Date/Time	Jan 1, 0001 A.D. to Dec 31, 9999 A.D. (precision to the nanosecond)
Varbyte	1 to 64,000 bytes	Binary	1 to 104,857,600 bytes
Varchar	1 to 64,000 bytes	String	1 to 104,857,600 characters

INDEX

D

- data object operation properties
 - input properties [22](#), [25](#)
 - target properties [24](#), [29](#)
- data object read operation
 - creating [31](#)
- data object write operation
 - creating [31](#)

M

- mapping run-time environments
 - Hive [33](#)
- mapping validation environments
 - Hive [33](#)

P

- PowerExchange for Teradata Parallel Transporter API
 - adapter overview [9](#)
 - run-time processing [35](#)
- PowerExchange for Teradata Parallel Transporter API configuration
 - database privileges [12](#)
 - JDBC settings [12](#)
 - prerequisites [10](#)
 - setting environment variables [11](#)

R

- read operation input properties
 - advanced [22](#)
 - sources [22](#)
- read operation target properties
 - column [24](#)
- row-level processing
 - duplicate rows [38](#)
 - extra rows [38](#)
 - missing rows [38](#)

T

- Teradata Connector for Hadoop
 - configuration file [14](#)
 - configuration options [13](#)
 - configuring [14](#)
 - enabling [15](#)
 - installing [14](#)
 - overview [13](#)
- Teradata data objects
 - importing [30](#)
 - overview [20](#)
 - overview properties [21](#)
 - sections [21](#)
- Teradata data types
 - mapping with transformation data types [42](#)
 - overview [42](#)
- Teradata mappings
 - example [34](#)
 - overview [33](#)
- Teradata Parallel Transporter connections
 - creating [18](#)
 - overview [16](#)
 - properties [16](#)
- Teradata PT API
 - error limit [36](#)
 - error tables [37](#)
 - log tables [37](#)
- Teradata PT API system operators
 - load [36](#)
 - stream [36](#)
 - update [36](#)

W

- write operation input properties
 - advanced [26](#)
 - column [25](#)
- write operation target properties
 - column [30](#)
 - keys properties [24](#), [30](#)